

LAW 2026

**The 20th Linguistic Annotation Workshop**

**Proceedings of the Workshop**

July 3, 2026

The LAW organizers gratefully acknowledge the support from the following sponsors.



©2026 Association for Computational Linguistics

Order copies of this and other ACL proceedings from:

Association for Computational Linguistics (ACL)  
317 Sidney Baker St. S  
Suite 400 - 134  
Kerrville, TX 78028  
USA  
Tel: +1-855-225-1962  
[acl@aclweb.org](mailto:acl@aclweb.org)

ISBN 979-8-89176-404-0

## Preface

Linguistic annotation of natural language corpora is the backbone of supervised methods of statistical natural language processing. The Linguistic Annotation Workshop (LAW) is the annual workshop of the ACL and ELRA Special Interest Group on Annotation (SIGANN), and it provides a forum for the presentation and discussion of innovative research on all aspects of linguistic annotation, including the creation and evaluation of annotation schemes, methods for automatic and manual annotation, use and evaluation of annotation software and frameworks, representation of linguistic data and annotations, semi-supervised “human in the loop” methods of annotation, crowd-sourcing approaches, and more. As in the past, the LAW will provide a forum for annotation researchers to work towards standardization, best practices, and interoperability of annotation information and software.

These proceedings include papers that were presented at the 20th Linguistic Annotation Workshop (LAW-XX), co-located with ACL 2026 in San Diego, California, United States, on July 3, 2026.

This edition of the workshop is the twentieth meeting of the ACL and ELRA Special Interest Group for Annotation. The first LAW workshop took place in 2007 at the ACL in Prague. Since then, the LAW has been held every year, consistently drawing substantial participation (both in terms of paper/poster submissions and participation in the actual workshop). This sustained engagement provides evidence that the LAW’s overall focus on annotation and data quality continues to be an important area of interest in the field, a substantial part of which relies on supervised learning from gold standard data sets and trustworthy evaluation in the era of large language models. This year, we received 41 direct submissions, out of which 20 papers have been accepted to be presented at the workshop, as long or short papers, or as posters.

In addition, LAW-XX features two invited talks by Marie-Catherine de Marneffe (UCLouvain) and Rachel Rudinger (University of Maryland), as well as a roundtable discussion that brings together diverse perspectives on open questions and future directions relevant to the workshop and its special theme.

The special theme of LAW-XX is *Errors in Annotation*. The special oral session highlights research on identifying and addressing annotation errors arising from objectively incorrect labels, annotations, or omissions, as distinct from cases of annotator disagreement. It also examines the role of error analysis in improving the quality of both human-annotated and LLM-generated datasets. As the field increasingly relies on large-scale annotated data, ensuring data quality has become a critical concern. Consequently, there is a growing need for methods, metrics, and tools that can systematically assess, quantify, and improve the quality of annotations across diverse data sources.

We extend our thanks to SIGANN for its financial support, and to the organizing committee for its continued dedication to the LAW workshops. We are also grateful to the ACL Workshop Office, Zeerak Talat, and the ACL 2026 Workshop Co-Chairs, Loic Barrault and Yang Zhao, for their guidance and support throughout the organization of the workshop. Above all, we would like to thank all the authors for submitting their papers to the workshop and our program committee members for their time, expertise, and thoughtful reviews, which made this workshop possible.

**The LAW-XX Program Co-Chairs:**  
Yang Janet Liu and Luke Gessler

# Organizing Committee

## Program Chairs

Yang Janet Liu, University of Pittsburgh  
Luke Gessler, Indiana University

## Organizing Committee

Amir Zeldes, Georgetown University  
Ines Rehbein, Mannheim University

## Keynote Talk

**Marie-Catherine de Marneffe**  
FNRS - UCLouvain - CENTAL  
2026-07-03 09:00:00 – Room: **La Jolla**

**Bio:** Marie-Catherine de Marneffe is a FNRS research associate and professor at UCLouvain. She obtained her PhD under the supervision of Chris Manning at Stanford University and worked 10 years in the Linguistics department at The Ohio State University as assistant then associate professor. Her main research interests are in computational pragmatics, building models that capture what people infer “between the lines”. Her 2012 dissertation highlighted the need to capture variation in human annotations. She is also one of the developers of the Universal Dependencies framework. Her research work has been funded by Google Inc., the National Science Foundation and the FNRS.

## **Keynote Talk**

**Rachel Rudinger**

University of Maryland

**2026-07-03 16:00:00 – Room: La Jolla**

**Bio:** Rachel Rudinger is an Assistant Professor of Computer Science at the University of Maryland. Her research interests lie in the areas of natural language understanding, commonsense reasoning, computational semantics/pragmatics, and issues of sociocultural fairness in NLP systems. She is a recipient of the National Science Foundation CAREER Award and holds a Ph.D. in Computer Science from Johns Hopkins University.

## Table of Contents

<i>Annotating Clinical Risk and Variation in Haitian Creole Medical Translation</i> Ludovic Mompelat, David Tézil and Rose Flaure Accilien .....	1
<i>Parser agreement and disagreement in L2 Korean UD: Implications for human-in-the-loop annotation</i> Hakyung Sung and Gyu-Ho Shin .....	12
<i>Rules-based system for Czech legal text readability</i> Kateřina Motalík Hodková, Ivan Kraus and Barbora Hladká .....	22
<i>Human-AI Annotation Error Auditing for Hebrew Diacritization with Frontier LLMs</i> Hillel Gershuni and Avi Shmidman .....	33
<i>Beyond Annotator Disagreement: Guideline-Induced Errors in Arabic Hate Speech Annotation</i> Wajdi Zaghouani .....	47
<i>When LLMs Disagree with Human Experts: Understanding LLM Annotation Failures in Nutrition Misinformation through Hierarchical Error Analysis using Seed Oil Narratives</i> Vishwaa Shah, Indika Kahanda and Andrea Arikawa .....	59
<i>Math-DB: A Discourse Framework for Mathematical Word Problems to Enhance LLM Reasoning</i> Mustafa Erolcan Er .....	75
<i>Cross-Linguistic Situation Entity Segmentation for Discourse Analysis in Diachronic English and German Text</i> Hanna Schmück, Veronika Urban, Xaver Krückl, Sonja Zeman, Claudia Claridge and Annemarie Friedrich .....	95
<i>UD-CHILDES-BG: a dependency treebank of Bulgarian child and child-directed speech</i> Mila Marcheva-Nash, Yasena Chantova, Tsvetina Kirilova, Ivelina Pavlova, Tsvetelina Stefanova, Yoana Vasileva and Weiwei Sun .....	113
<i>IndiAnn: A Web-based Annotation Platform for Indic Languages</i> Bandaru Lavadeep, Ritwik Raghav and Abhik Jana .....	130
<i>Designing Annotation Guidelines for Trait-Based Arabic Automated Essay Scoring: A Systematic Methodology</i> Walid Massoud, Houda Bouamor, Abdelrahman Abdel Latif Hussein and Abdullah Mohamed Mohamed Zekri .....	146
<i>Revisiting Faithfulness Annotations for Long-form Summaries</i> Yang Zhong, Yang Janet Liu and Diane Litman .....	158
<i>Completing and Validating the Re-Aligned Switchboard Dialog Act Corpus</i> Run Chen, Zihao Tao, John Prado, Ignazio LaManna, Ryan Puterbaugh, Mim Datta and Julia Hirschberg .....	173
<i>Not Worth Mentioning? A Pilot Study on Salient Proposition Annotation</i> Amir Zeldes, Katherine Conhaim and Lauren Levine .....	178
<i>LLMs as annotators of credibility assessment in Danish asylum decisions: evaluating classification performance and errors beyond aggregated metrics</i> Galadrielle Humblot-Renaux, Mohammad N. S. Jahromi, Rohat Bakuri-Jørgensen, Marieke Anne Heyl, Asta S. Stage Jarlner, Maria Vlachou, Anna Murphy Høgenhaug, Desmond Elliott, Thomas Gammeltoft-Hansen and Thomas B. Moeslund .....	187

<i>Cracks in the Bridge—or A Bridge Too Far? Comparing Human and LLM Errors in the Annotation of Bridging Anaphora</i>	
Lauren Levine and Amir Zeldes .....	219
<i>Clustering Analysis for Error Detection in Named Entity Recognition Datasets</i>	
Matthew Flynn, Timothy Obiso, Sam Newman and Constantine Lignos .....	229
<i>When Ground Truth Disagrees: A Human-in-the-Loop Audit of Annotation Errors in High-Stakes Crash Narratives</i>	
Md Sajjad Hossain, Lin Li, Judy A. Perkins, John Clary and Joel Meyer .....	241
<i>Prompts in the Wild: A Large Analyzed Collection of Transactional Prompts in Code</i>	
Victoria Basmov, Yoav Goldberg and Reut Tsarfaty .....	257
<i>TalkTag: Fine-Grained Morphosyntactic Error Annotation for Transcribed Speech</i>	
Shamira Venturini, Oliver Hennhöfer, Steffen Kinkel and Jannik Strötgen .....	309

# Program

**Friday, July 3, 2026**

08:45 - 09:00 *Opening Remarks*

09:00 - 09:45 *Keynote 1: Marie-Catherine de Marneffe*

09:45 - 10:30 *Talks*

*Human-AI Annotation Error Auditing for Hebrew Diacritization with Frontier LLMs*

Hillel Gershuni and Avi Shmidman

*LLMs as annotators of credibility assessment in Danish asylum decisions: evaluating classification performance and errors beyond aggregated metrics*

Galadrielle Humblot-Renaux, Mohammad N. S. Jahromi, Rohat Bakuri-Jørgensen, Marieke Anne Heyl, Asta S. Stage Jarlner, Maria Vlachou, Anna Murphy Høgenhaug, Desmond Elliott, Thomas Gammeltoft-Hansen and Thomas B. Moeslund

*Revisiting Faithfulness Annotations for Long-form Summaries*

Yang Zhong, Yang Janet Liu and Diane Litman

10:30 - 11:15 *Coffee Break*

11:15 - 12:00 *Talks*

*UD-CHILDES-BG: a dependency treebank of Bulgarian child and child-directed speech*

Mila Marcheva-Nash, Yasena Chantova, Tsvetina Kirilova, Ivelina Pavlova, Tsvetelina Stefanova, Yoana Vasileva and Weiwei Sun

*When Ground Truth Disagrees: A Human-in-the-Loop Audit of Annotation Errors in High-Stakes Crash Narratives*

Md Sajjad Hossain, Lin Li, Judy A. Perkins, John Clary and Joel Meyer

*Rules-based system for Czech legal text readability*

Kateřina Motalík Hodková, Ivan Kraus and Barbora Hladká

12:00 - 13:30 *Lunch*

13:30 - 14:15 *Talks*

Friday, July 3, 2026 (continued)

*TalkTag: Fine-Grained Morphosyntactic Error Annotation for Transcribed Speech*

Shamira Venturini, Oliver Hennhöfer, Steffen Kinkel and Jannik Strötgen

*Semantic-pragmatic Annotations in the Prague Dependency Treebank*

Marie Mikulová, Eva Hajičová, Jiří Mírovský, Anna Nedoluzhko, Michal Novák, Pavlína Synková, Jan Štěpánek, Barbora Štěpánková and Jan Hajič

*Not Worth Mentioning? A Pilot Study on Salient Proposition Annotation*

Amir Zeldes, Katherine Conhaim and Lauren Levine

14:15 - 14:30 *Short Break*

14:30 - 15:30 *Posters*

*Designing Annotation Guidelines for Trait-Based Arabic Automated Essay Scoring: A Systematic Methodology*

Walid Massoud, Houda Bouamor, Abdelrahman Abdel Latif Hussein and Abdullah Mohamed Mohamed Zekri

*Cracks in the Bridge—or A Bridge Too Far? Comparing Human and LLM Errors in the Annotation of Bridging Anaphora*

Lauren Levine and Amir Zeldes

*Prompts in the Wild: A Large Analyzed Collection of Transactional Prompts in Code*

Victoria Basmov, Yoav Goldberg and Reut Tsarfaty

*When LLMs Disagree with Human Experts: Understanding LLM Annotation Failures in Nutrition Misinformation through Hierarchical Error Analysis using Seed Oil Narratives*

Vishwaa Shah, Indika Kahanda and Andrea Arikawa

*Math-DB: A Discourse Framework for Mathematical Word Problems to Enhance LLM Reasoning*

Mustafa Erolcan Er

*Parser agreement and disagreement in L2 Korean UD: Implications for human-in-the-loop annotation*

Hakyung Sung and Gyu-Ho Shin

*Completing and Validating the Re-Aligned Switchboard Dialog Act Corpus*

Run Chen, Zihao Tao, John Prado, Ignazio LaManna, Ryan Puterbaugh, Mim Datta and Julia Hirschberg

**Friday, July 3, 2026 (continued)**

*Clustering Analysis for Error Detection in Named Entity Recognition Datasets*

Matthew Flynn, Timothy Obiso, Sam Newman and Constantine Lignos

*Annotating Clinical Risk and Variation in Haitian Creole Medical Translation*

Ludovic Mompelat, David Tézil and Rose Flaure Accilien

*EVADE: LLM-Based Explanation Generation and Validation for Error Detection in NLI*

Longfei Zuo, Barbara Plank and Siyao Peng

*IndiAnn: A Web-based Annotation Platform for Indic Languages*

Bandaru Lavadeep, Ritwik Raghav and Abhik Jana

*Beyond Annotator Disagreement: Guideline-Induced Errors in Arabic Hate Speech Annotation*

Wajdi Zaghouni

*Cross-Linguistic Situation Entity Segmentation for Discourse Analysis in Diachronic English and German Text*

Hanna Schmück, Veronika Urban, Xaver Krückl, Sonja Zeman, Claudia Claridge and Annemarie Friedrich

15:30 - 16:00 *Coffee Break*

16:00 - 16:45 *Keynote 2: Rachel Rudinger*

16:45 - 17:15 *Roundtable Discussion: Open Questions & Future Directions*

17:15 - 17:30 *Closing Remarks*

# Annotating Clinical Risk and Variation in Haitian Creole Medical Translation

**Ludovic Mompelat**  
University of Miami  
Miami, USA  
ludovic.mompelat@miami.edu

**David Tézil**  
University of Alabama  
Tuscaloosa, USA  
dtezil@ua.edu

**Rose Flaure Accilien**  
University of Miami  
Miami, USA  
rxa1262@miami.edu

## Abstract

We present an annotation schema for Haitian Creole medical translation that makes clinical risk and sociolinguistic variation explicit while remaining lightweight enough for small expert teams. The schema includes binary fields for overall acceptability, severity of potential misunderstanding, and foreign-influence cues, along with conditional error tags aligned with Multidimensional Quality Metrics (MQM), commonly used in the medical domain, for interoperability. Through three rounds of annotation and adjudication we achieve stable inter-annotator agreement and release a gold dataset of 152 EN→HC medical sentence pairs. A simple classifier–labeller baseline demonstrates that acceptability and severity are reliably learnable under data scarcity, while foreign-influence judgments remain limited by prevalence. These results show that clinically oriented, variety-sensitive annotation can both support immediate screening of patient-facing translations and provide reward-ready signals for future preference-based MT and LLM fine-tuning.

## 1 Introduction

Despite growing efforts in low-resource machine translation, and especially in medical contexts, Haitian Creole (HC) remains critically underrepresented (Mompelat, 2025). This further limits the availability and development of effective translation tools capable of adequately supporting the efforts of human interpreters and medical providers. In regions like Miami, which is a diasporic hub for HC speakers, healthcare disparities are exacerbated by language barriers, with existing translation systems failing to address intra-language variation, code-switching, and the absence of standardized medical terminology. The stakes in effective and adequate medical communication are particularly high since lexical choices and grammatical structures can directly affect patient comprehension,

trust, adherence, and safety. Language barriers in healthcare have also been shown to reduce patient and provider satisfaction, compromise care quality, and increase miscommunication and associated costs (Pérez-Escamilla et al., 2010; Al Shamsi et al., 2020). In the particular case of HC, these risks are compounded by intra-linguistic variation across basilectal (Creole-exclusive) and mesolectal (French-influenced) varieties (Bickerton, 1973), and inter-linguistic contact with French, English and Spanish, particularly in Miami. Available HC medical data are uneven in format, size, and quality, making them difficult to aggregate or use directly for modeling. To address this, we began by normalizing and preprocessing the CMU Haitian Creole medical corpus<sup>1</sup>, which revealed the need for a structured annotation framework. The schema we introduce in this paper guides both dataset cleaning and evaluation, capturing the sociolinguistic and clinical dimensions that matter most for safe and effective HC medical communication. This pilot is designed as a schema validation study under low-resource conditions and intended for interdisciplinary annotation efforts; simple enough for linguists, but precise and accurate enough to align with domain-specific standards of annotation such as Multidimensional Quality Metrics (MQM) used in the medical context. Annotation reliability and label learnability are therefore the primary objectives. To check practical utility, we train a small, interpretable classifier–labeller on the adjudicated EN→HC pairs with five-fold cross-validation. The model outputs a probability for each binary field and suggests error tags. In this setup, we find that (i) deciding whether a translation needs improvement is reliable; (ii) error detection is tractable once quality issues are identified but requires more instances of rarer terminology and orthographic errors; and (iii) detecting foreign-influence phe-

<sup>1</sup><http://www.speech.cs.cmu.edu/haitian/>

nomena remains data-limited at current coverage. Operationally, the *Severity* probabilities let reviewers look first at the most consequential cases, and the pre-filled error tags speed up adjudication.

## 2 Background and Motivation

### 2.1 Clinical Stakes of Language Choice and Access in Healthcare

When patients and providers do not share a language, communication breaks down and safety suffers. Recent reviews synthesize a consistent picture: language barriers drive misunderstandings, reduce satisfaction on both sides, and are linked to preventable errors and worse outcomes (Al Shamsi et al., 2020).

Professional interpreters remain the standard of care bridging the gap (often virtually) between patient and healthcare practitioners. Their use is associated with clearer communication, guideline-concordant treatment, and better outcomes (Karlner et al., 2007; Jacobs et al., 2004). Yet the availability of qualified interpreters is not guaranteed, particularly for low-resource languages where on-demand access is limited and scheduling delays are common. In these settings, clinicians often rely on ad-hoc interpreting, a practice shown to increase error rates and miscommunication, especially in pediatrics and emergency care; by contrast error rates fall significantly with trained interpreters (Flores et al., 2003, 2012; Divi et al., 2007).

Within that reality, technology often fills the gaps. A scoping review by Kreienbrinck et al. (2025) distinguishes two main tool types: fixed-phrase systems and Machine Translation apps. Phrasebooks and menu-driven tools can support brief, predictable exchanges but fail outside their limited inventories and require constant maintenance for domain coverage (Hudelson and Chappuis, 2024; Noack et al., 2021; Spechbach et al., 2019). MT systems, in turn, can handle arbitrary content and more languages, but accuracy remains inconsistent; studies commonly warn against unsupervised use in safety-critical communication, particularly beyond high-resource languages (Panayiotou et al., 2019; Halimi and Bouillon, 2019; Hwang et al., 2022). The consensus emerging from both research and clinical guidance is therefore clear: technology can mitigate access gaps when no interpreter is available, but it must operate under supervision and never replace qualified professionals.

Policy frameworks reinforce this balance between access and safety. The U.S. Department of Health and Human Services' National CLAS Standards mandate free language assistance, proactive notification of its availability, and interpreter competence (U.S. Department of Health and Human Services, 2013). The NIH Clear Communication Initiative and the CDC's "Everyday Words for Public Health" extend these principles to written materials, prescribing plain-language phrasing and comprehension testing with target audiences (National Institutes of Health, 2025; Centers for Disease Control and Prevention, 2016). Similarly, the World Health Organization emphasizes accessibility so that non-experts can act on health information (World Health Organization, 2021), and Translators without Borders (CLEAR Global) operationalize these ideals through multilingual glossaries and tools for crisis response (Translators without Borders, 2020). Together, these policies frame linguistic equity as both a legal and ethical duty rather than a technical convenience.

HC makes these tensions concrete. It exemplifies a language where interpreter availability is limited and technological coverage is uneven, forcing practitioners to rely on improvised, inconsistent solutions. In such settings, variety-appropriate phrasing becomes a matter of safety as well as clarity: what matters is not only that a translation is correct, but that it is intelligible and culturally legible to monolingual HC speakers who cannot fall back on French.

### 2.2 Haitian Creole and Medical Communication

HC is the native language of almost the entire population of Haiti; approximately 90–95% of Haitians speak it fluently, and for many, it is their only language (Dejean, 2010; Hebblethwaite, 2012). Globally, HC counts around 10–12 million speakers, making it the most widely spoken Creole (Valdman et al., 2017). In the United States, Haitian Americans number over 1.2 million, with substantial communities in South Florida (especially Miami and Orlando), New York City, and Boston (US Census 2024). In Florida specifically, HC ranks as the third most spoken language after English and Spanish, a fact reflected in local service provision and public communications.

Focusing on the city of Miami, one of the many problems faced by the HC community is access to appropriate care, and while many factors and vari-

ables come into play to explain the discrepancies in quality of and access to care, such as cultural literacy from the medical staff (Campbell, 2012), or financial/employment difficulties faced by the local community (in Miami, for example, and especially in Little Haiti) (Ryan et al., 2004; Kobetz et al., 2009, 2010; Menard et al., 2010), we are particularly interested in tackling the language barrier factor. Public agencies and health systems routinely provide language assistance and translated materials—including in HC—so that patients can access basic information and services. This is visible at the federal level (e.g., CDC language-assistance pages and multilingual health resources) and locally in Miami, where city and county plans specify notices and summaries in English, Spanish, and HC. In everyday practice, however, HC-specific tools remain patchy, so clinicians and interpreters often lean on improvised mixes of translation apps and glossaries. Additionally, during emergencies, agencies and partners have even circulated outdated or inaccurate purpose-built HC materials (e.g., cholera and earthquake response guides, triage glossaries), underscoring the need for up-to-date, domain resources. Finally, in the NLP space, HC MT is often routed via French as a pivot language, a standard low-resource strategy which at times exacerbates existing gaps in direct HC support or create new translation issues (Dholakia and Sarkar, 2014); major commercial systems like DeepL also do not list HC among supported languages.

### 2.3 Haitian Creole Linguistic Variation and Gaps in Current Models

HC exhibits systematic sociolinguistic variation often described along a continuum between *kreyòl swa* (prestige variety shaped by French contact) and *kreyòl rèk* (basilectal, Creole-like, monolingual variety) (Tezil, 2022; Valdman, 2015; DeGraff, 2005). This structural and lexical diversity matters particularly in the medical domain. For example, terms like *dyabèt* (from French *diabète*) are readily legible to bilinguals, but monolingual speakers may instead expect variants such as *maladi sik* (‘sugar disease’). The linguistic continuum in HC yields multiple plausible translations for the same clinical concept, but with very different implications for comprehension and patient trust. For medication, *medikaman* suggests biomedical precision, whereas *remèd/remmèd* evokes a broader category that includes more-culturally tied folk or herbal remedies, requiring careful contextualization. The

wrong lexical choice can introduce confusion or mistrust, even when the translation is literally accurate. (Valdman et al., 2017; DeGraff, 2005)

Other contrasts show the risks of outright misunderstanding. A provider’s note about chest pain may be translated with *pwatrin* (chest) or *lestomak* (stomach), leading to clinically dangerous miscommunication. In substance-use screening, asking a patient if they are a *tafyatè* (literally ‘rum/alcohol drinker’) to ask if they drink alcohol (i.e. regularly, occasionally, or not) risks offense: the term is widely understood to mean ‘alcoholic,’ carrying stigma that can prevent disclosure. These examples underscore that in HC medical contexts, translation accuracy cannot be reduced to word-for-word fidelity. Audience design—deciding which variety and phrasing suit the patient population—becomes a clinical and critical safety issue. Our schema takes these linguistically-based observations into account by requiring annotators to score both correctness and record possible variety-specific variants of lexical items or phrases, producing potential signals, useful for language modeling, that surface when outputs are comprehensible but inappropriate for their intended audience.

Beyond variation, outdated lexicographic resources introduce practical constraints. Widely used dictionaries and phrasebooks predate post-2010 usage (e.g., not including the earthquake-era neologism *goudougoudou* for trauma/distress), and some entries reflect mesolectal/French bias or hybridized forms that are not transparent to younger or monolingual users; many sources also remain non-digital, complicating updates and integration into NLP pipelines (Freeman, 1997; Heuretélou et al., 2000).

Current MT systems for HC rarely capture the internal diversity of the language. They might default toward French-influenced lexical and syntactic patterns, which can alienate or confuse monolingual readers who are less familiar with acrolectal forms. Yet the inverse is also true: many “standard” basilectal forms used in official or corpus materials do not necessarily reflect how either monolingual or bilingual speakers actually speak in clinical or everyday settings. In practice, preferred usage often lies somewhere between the two poles. What counts as “plain” or “appropriate” HC depends on audience, region, and communicative setting. This fluidity makes it risky to treat HC as a homogeneous variety or to assume that basilectal vocabulary is always the most accessible (Lewis, 2010;

Mompelat, 2025).

These tensions are reflected in our annotated data, where fluency problems frequently co-occur with *foreign-influence* cues such as overly Frenchified function words, calques, or hybridized syntax—evidence that linguistic quality and sociolinguistic variation are intertwined. Our schema is designed to make this complexity explicit rather than to enforce a single normative variety. A sentence can display *foreign influence* without being wrong: in those cases, annotators may propose a basilectal variant when the translation is accurate but mesolectal or French-influenced, allowing for variety-sensitive alternatives. When *foreign influence* co-occurs with a *quality* issue, however, the influence often signals a deeper translation problem (e.g., forms that sound unnatural or implausible to any HC speaker).

### 3 Annotation Schema

We design a lightweight schema aimed at maximizing reliability in small expert teams while retaining the expressive power needed to capture three aspects central to HC medical communication: sensitivity to variety and contact-driven variety differences, explicit attention to clinical risk, and a minimal set of error tags aligned with widely used translation-quality taxonomies for downstream interoperability. Existing frameworks such as the Multidimensional Quality Metrics (MQM) (Lommel et al., 2014) or the NCC MERP taxonomy (National Coordinating Council for Medication Error Reporting and Prevention (NCC MERP), 2022) do not explicitly encode clinical risk and intralanguage variation, which are central in HC medical communication; our schema extends both minimally to capture these dimensions while preserving interoperability.

#### 3.1 Guiding principles and rater instructions

Annotators worked from compact definitions, examples, and decision tests that privileged consistency over granularity. *Foreign influence* was judged by comparing outputs to reputable lexicographic sources (Valdman, 2015; Valdman et al., 2017) and by flagging recognizable contact-induced patterns in lexicon or morphosyntax (e.g., French-leaning terms, calques, complementizer *ke* in acrolectal styles). *Quality* was defined according to monolingual and bilingual HC speakers acceptability, grammaticality and understandability

criteria. *Severity* was introduced as a separate, risk-oriented lens, grounded in clinical stakes rather than linguistic well-formedness (Flores et al., 2003, 2012).

Where a sentence was judged improvable (quality = 1), annotators had to provide a corrected version and could propose additional variants to the corrected forms that are as suitable depending on the audience. Raters also noted uncertainty, which we used to flag items for adjudication. Instructions were refined across rounds as systematic sources of disagreement emerged, with clarifications added for recurring borderline cases.

#### 3.2 Labels and decision criteria

*Foreign influence*: binary field marking whether a HC output leans toward French- or English-influenced, or otherwise non-Creole, forms in its lexical choices or morphosyntax. The rationale is both sociolinguistic and diagnostic. On the one hand, HC exhibits socially conditioned contrasts often described along the *kreyòl swa-kreyòl rèk* continuum (Fattier-Thomas, 1984; Tézil, 2024), where the choice in the linguistic variety influences intelligibility and trust in patient-facing materials. On the other hand, *foreign influence* also captures translation-induced interference from source languages such as English negative polarity item behavior and negation structure as shown in example (1) and which corrected sentence is the negative concord sentence “men ou *pa* konnen *anyen ankò*”.

- (1) Men ou konnen okenn lòt bagay  
but you know none other thing  
'But you know nothing else.'

The label therefore serves a dual function: to flag variety alignment issues that may affect accessibility, and to detect morphosyntactic intrusions introduced by translators or machine systems.

*Severity*: binary risk label indicating whether the content of a sentence, if misunderstood or mistranslated, could plausibly alter a patient action or clinical decision (e.g., dosing, contraindications, symptom actionability, discharge instructions).

We draw on the NCC MERP Index, which organizes medication errors along a spectrum from potential risk (Category A) to patient harm and death (Categories E–I), but we do not attempt to reproduce these fine-grained distinctions (National Coordinating Council for Medication Error Reporting and Prevention (NCC MERP), 2022). Instead,

we collapse this spectrum into a binary decision: sentences are marked as high-risk (*Severity*=1) if they involve information whose misinterpretation could have clinical consequences, regardless of whether that consequence would be minor or severe, and low/no risk (*Severity*=0) otherwise.

Crucially, this label does not assess whether a translation is erroneous, but whether the underlying *information type* carries potential clinical risk under miscommunication. We use, instead, the label *Error - terminology* for cases of clear erroneous domain-specific terminology. This design separates linguistic correctness from risk exposure and enables straightforward triage of patient-facing materials.

*Quality*: binary acceptability decision, clear and accurate for the intended audience vs. improvable. This criterion aligns with the clinical importance of plain-language communication and cultural fit. When a sentence is judged as needing improvement (i.e., *Quality* = 1), raters apply minimal *error tags* to diagnose the problem in MQM-compatible terms (Lommel et al., 2014): (i) terminology or accuracy problems, (ii) fluency or adequacy breakdown, and (iii) orthographic errors.

Example (2) is the proposed translation of a question on *casual* alcohol consumption illustrates how these dimensions interact but remain distinct. The translation uses *tafyatè* ‘alcoholic’, which is semantically inappropriate and socially negatively marked (i.e. someone diagnosed as suffering from alcoholism) in this context. This triggers *Quality* = 1 and a *terminology* error. At the same time, the sentence is marked *Severity* = 1 because misunderstanding or misframing alcohol use can affect patient disclosure and downstream care decisions. The example shows that *severity* is not reducible to linguistic error: it reflects the potential clinical consequences of miscommunication, even when the sentence is otherwise fluent.

(2) èske ou se yon tafyatè  
 Q 2SG COP DET alcoholic  
 ‘Are you an alcoholic?’ (intended: ‘Do you drink alcohol?’)

Restricting error tags to improvable sentences reduces noise and mirrors clinical review workflows, where reviewers first decide whether the material is usable, and only then identify what needs fixing.

Finally, because multiple semantically equivalent lexical items and phrasings are possible and

variety- or region-dependent in HC, we also record *variants*: a binary flag indicating whether a clearly viable alternative exists that preserves meaning but may better match the patient’s idiolect (e.g., *bwason alkolize*, *tafya*, or *kleren* ‘alcoholic drink’).

In section 5, we apply the schema to a small, adjudicated EN→HC medical set and probe its modeling potential with a baseline classifier–labeller. Results show that acceptability detection is robust, error typing is tractable when quality issues are present, and foreign-influence detection is limited more by data coverage than by label design.

## 4 Annotation Workflow

### 4.1 Rounds and Sampling

Annotators were two bilingual Haitian Creole–English professionals: one certified medical interpreter with clinical experience in South Florida, and one trained linguist specializing in Haitian Creole variation and translation. Annotation was conducted in three successive rounds designed to iteratively refine the schema and improve inter-annotator agreement:

- **Phase 1:** 300 sentences were independently annotated by two annotators using the initial guidelines.
- **Phase 2:** 100 new sentences were introduced, with adjustments to field definitions and decision rules based on Phase 1 disagreements.
- **Phase 3:** 52 sentences from Phase 2 that were re-annotated after updating the schema, and 50 additional sentences were newly annotated to target categories that remained problematic, enabling focused reliability testing.

Adjudication followed each round. Disagreements and low-confidence cases were reviewed with reference to guidelines, lexicons, and audience assumptions. Recurring edge cases triggered refinements to instructions, reducing drift. The final round sampled additional material with an eye toward stress-testing categories that remained difficult such as foreign-influence phenomena and rare error types. All items were presented with the English source sentence, the HC hypothesis, and minimal discourse context. The pilot experiment later in the paper uses the 152 EN→HC medical pairs drawn from across Phases 2 and 3, after normalization and adjudication. The next annotation effort

will consist in re-annotating the 300 sentences from Phase 1.

## 4.2 Agreement and Adjudication

Reliability was assessed at the field level using Cohen’s  $\kappa$  on raw labels prior to adjudication and Krippendorff’s  $\alpha$  for nominal data as a robustness check (Artstein and Poesio, 2008). Agreement is reported separately for each binary label; for error tags, we restrict computation to the subset where quality was judged improvable, consistent with the schema’s conditional design.

Agreement improved over rounds. In Phase 2 we observe substantial reliability for severity ( $\kappa = 0.709$ ) but only moderate to fair scores for quality (0.454), fluency (0.409), and foreign influence (0.251). By Phase 3, terminology rises sharply ( $\kappa = 0.850$ ) after dictionary-backed clarifications, while quality (0.480) and typo (0.490) improve modestly. Severity declines somewhat (0.561) because Phase 3 targeted more ambiguous, higher-stakes content. Foreign influence remains difficult (0.203), reflecting both low prevalence and fuzzy sociolinguistic boundaries.

From the final adjudication pass, we compile a gold set with fully resolved labels across all schema dimensions. Disagreements were resolved collaboratively through guideline clarification and consensus review, ensuring that the dataset reflects adjudicated, reproducible judgments. The dataset is structured to support both sentence-level evaluation and span-based processing: corrected translations and audience-appropriate variants are stored as explicit token-level mappings, enabling fine-grained alignment and downstream use in sequence-to-sequence or correction modeling.

## 4.3 Toward Reward Modeling for Haitian Creole MT

Recent work in preference learning and reinforcement learning from human feedback (RLHF) shows how annotated judgments can steer large language models toward outputs that reflect human preferences and contextual appropriateness rather than surface likelihood alone (Wang et al., 2024; Dong et al., 2024; Ouyang et al., 2022). Within MT, reward modeling has been used to rerank candidate translations and fine-tune generation toward domain-specific criteria such as fluency, adequacy, and user-centered preferences (Kreutzer et al., 2018; Lyu et al., 2023).

The adjudicated labels—*quality* (acceptability),

*error type* when problems occur, *foreign influence* as a signal of language contact influence and linguistic variation, and *severity* as a clinical-risk indicator—map naturally onto preference data and potential reward functions. In this framing, severity can serve as a cost-sensitive weight, quality as a binary accept/reject signal, and foreign influence as a dimension of idiolect appropriateness. Together, these labels provide the basis for both pairwise preference modeling and scalar reward training, and can be scaled through active learning once classifiers flag likely error cases.

In the pilot study that follows, we use the schema purely in a supervised setup with a simple classifier–labeller baseline. Results show that acceptability detection is reliable, error typing is tractable once quality issues are flagged, and foreign-influence judgments remain limited chiefly by data coverage. Future work will extend these same labels into reward modeling pipelines that bias generation toward patient-facing plain language and culturally appropriate renderings in HC medical communication.

## 5 Pilot Study

### 5.1 Setup and model

The baseline is intentionally simple and interpretable, serving as a lower bound on label learnability (Table 1).

Table 1: Classifier setup summary.

Component	Description
Features	TF–IDF (EN+HC), length, diacritics, overlap
Models	Logistic regression (per label), OVR for tags
Evaluation	5-fold Cross-validation, stratified
Outputs	Probabilities + binary labels

Each EN–HC pair is represented by a shared feature vector combining TF–IDF  $n$ -grams and lightweight numeric cues. We train one logistic-regression classifier per binary label (*quality*, *severity*, *foreign influence*, *variants*) and use a one-vs-rest setup for multi-label error tags. Evaluation uses five-fold cross-validation with stratification on the joint (*severity*, *quality*) label.

Feature design targets HC-specific signals while remaining interpretable: orthographic patterns (diacritics), length-based cues, and lexical overlap across source, hypothesis, correction, and variant capture alignment and variation without relying on contextual embeddings, which we leave for future

work.

## 5.2 Results

We compare against a prevalence-matched random baseline, which provides a lower bound under class imbalance. We report precision, recall, and F1 for each binary task, alongside the prevalence of positives (e.g., 18/152 for *foreign influence*).

Across labels, our model consistently exceeds this baseline, indicating that the annotation signals are recoverable rather than driven by label distribution. Because *severity* is intended to triage review, we compute a “limited-review recall”: sentence pairs are ranked by their predicted probability of being high-risk, and we simulate reviewing only the top 10%, 20%, or 30% of that ranked list. The metric reports the proportion of all truly high-risk items recovered within each reviewed portion. This measures how well model scores prioritize scarce human attention. Finally, we compute macro-F1 and per-tag F1 for *fluency*, *terminology*, and *typo* on that subset.

Table 2 summarizes performance for the four binary labels. *Quality* is detected with high reliability (F1 = 0.94), and *variants* are captured well (F1 = 0.89). Performance on *severity* is solid (F1 = 0.77), above the random baseline (0.66), while *foreign influence* remains more limited (F1 = 0.36), reflecting both low prevalence (11.8%) and the difficulty of separating contact-induced forms from HC-internal variation.

For *severity*, probability outputs support ranking. Because high-risk cases represent 66.4% of the dataset, randomly inspecting 10% of sentences would recover about 10% of them. The model instead recovers 14.9%, rising to 28.7% and 42.6% at 20% and 30% inspection budgets, corresponding to a 1.42–1.49× improvement over random. High-risk cases are therefore concentrated near the top of the ranking, making the score effective for prioritization.

Error typing (Table 3) shows a clear dependence on label prevalence. *Fluency* errors are detected reliably (F1 = 0.85), while *terminology* is moderate (F1 = 0.35) and *typo* remains weak (F1 = 0.17), consistent with their relative scarcity. This pattern suggests that the model captures frequent structural issues more readily than rarer, lexically specific phenomena.

Variance across folds is low for the main tasks (e.g., *quality*:  $0.941 \pm 0.023$ ; *severity*:  $0.767 \pm 0.044$ ), indicating stable estimates. Higher vari-

Table 2: Binary Classification Tasks: positive counts and P/R/F1 (5-fold out-of-fold,  $n=152$ ).

Label	Positives (#)	Prec.	Rec.	F1
Quality	53	0.980	0.906	0.941
Severity	101	0.837	0.713	0.770
Foreign infl.	18	0.296	0.444	0.356
Variants	54	0.920	0.852	0.885

Table 3: Error typing on the *quality=1* subset ( $n=53$ ): per-tag P/R/F1 and prevalence.

Tag	Positives (#)	Prec.	Rec.	F1
Fluency	45	0.864	0.844	0.854
Terminology	9	0.375	0.333	0.353
Typo	5	0.143	0.200	0.167
Macro-F1				0.458

ance for *foreign influence* ( $0.353 \pm 0.216$ ) reflects class imbalance and the limited number of positive instances.

## 6 Discussion

### 6.1 Linking Annotation, Sociolinguistics, and the Pilot

The schema was designed around two constraints distinctive to HC medical communication: contact-driven linguistic variation and clinical risk. The annotation rounds operationalized these with compact binary fields and conditional error tags, and the inter-annotator agreement profile revealed which judgments were inherently stable versus which demanded sharper guidance. The pilot then tested how learnable these labels are under data scarcity. Together, the findings line up: categories that humans applied consistently also supported robust automatic detection, while low-prevalence or conceptually fine-grained categories remain the main bottleneck.

Qualitatively, errors cluster into three recurring patterns. First, lexical misselection, including stigmatizing or overly specific terms (e.g., *tafyatè* ‘alcoholic’ instead of a neutral phrasing such as *èske w konn bwè* ‘do you drink?’); these cases are typically captured under terminology errors but may also carry clinical risk. Second, morphosyntactic interference, often reflecting transfer from English or French (e.g., calqued terminology ‘chest’ *#pwatrin* vs. *lestomak* or compositional phrasing such as ‘all three’ *\*tout twa* vs. *touletwa*), which the model tends to capture under fluency. Third, register and variety mismatches, where otherwise

correct translations rely on forms that are not uniformly shared across speakers (e.g., 'how are you feeling' *èske sa va* vs. *èske ou anfòm*, or 'alcoholic beverage' *bwason alkolize* vs. *bwe lalkol*), frequently surfacing in the foreign-influence and variants fields. These patterns align with the quantitative results: frequent structural issues (fluency) are reliably detected, while lexically specific or contact-sensitive distinctions (terminology, foreign influence) remain more variable under current data coverage.

*Quality* is the clearest case. Moving to a binary acceptability decision improved rater agreement, and the baseline model achieved high F1 on the same task. This alignment is expected: the features (TF-IDF and simple text-shape cues) capture adequacy and fluency in exactly the way the rubric defines them.

*Severity* shows a complementary profile. Annotation agreement dipped on ambiguous, high-stakes content such as dosing or discharge instructions, reflecting real borderline cases. Yet the model's probability scores effectively ranked these cases: limited-review recall showed that inspecting only a small slice of the highest-probability items recovered a large share of true high-risk instances. This is well suited to clinical triage, where the goal is to prioritize review rather than to enforce a single threshold.

*Foreign influence* remains the most challenging dimension. Performance is modest (F1 = 0.356) and highly variable across folds ( $\pm 0.216$ ), reflecting both label scarcity (11.8%) and conceptual subtlety. This instability points to a data limitation rather than a flaw in the schema: expanding coverage with targeted examples will be necessary to stabilize this signal. Because French is HC's lexifier, many French-origin elements are part of the language's core system and appear even in basilectal varieties. Yet contemporary contact with French continues to shape new mesolectal forms that are not shared by all speakers.

*Error tags* show a similar division. Terminology agreement improved once definitions were tightened and dictionary attestations used, while fluency emerged as both common and learnable. Terminology and orthographic errors remain rare but matter precisely because they can undermine safety and credibility.

Finally, *variants* demonstrate the schema's core sociolinguistic insight: multiple renderings may be "correct" but not equally appropriate for monolin-

gual versus bilingual audiences. Both annotators and the model handled this field well, and it provides a natural bridge to preference-based modeling.

## 6.2 From Pilot to Practice: Applications and Next Steps

The pilot turns the schema into a usable signal. Even with a compact model, the labels support screening of HC medical translations before human review. In practice, predicted *quality* flags likely problematic sentences, while *severity* probabilities rank them so that the riskiest material appears first. The limited-review analysis shows that this ordering makes scarce human attention more effective—a good match for clinical and public-health settings.

These same scores also help grow the dataset where it matters. Because error typing is meaningful once *quality*=1, the model can pre-label likely problems and defer to annotators where its confidence is low. This is classic active learning: propose, adjudicate, and retrain. *Foreign influence* is the clearest target. Our pilot quantified its scarcity, and now we can seek out additional examples from both Haiti and diaspora contexts, anchored in lexicographic attestations.

For evaluation and system steering, the schema becomes an actionable framework:

1. Require *quality*=0 for release,
2. Use *severity* to gate additional review,
3. Treat *variants* as a preference space for reranking toward monolingual or bilingual audiences,
4. Map error tags directly to edits—terminology errors suggest glossary fixes, fluency errors prompt rewrites, and typos trigger orthographic normalization.

As the dataset scales, these labels are reward-ready: they can drive pairwise preference modeling or scalar reward functions that bias MT toward plain, audience-appropriate HC while down-weighting foreign-influence intrusions.

A key extension will be the integration of curated glossaries and span-level mappings, already preserved in the normalized CSV. Grounding model suggestions in attested HC terminology rather than unconstrained paraphrase will make outputs more reliable and interpretable. In the longer term, the

same infrastructure can support dynamic adaptation—systems that learn to select variants aligned with a patient’s idiolect or speech profile, bridging linguistic diversity and clinical safety in real time.

In short, the experiment shows that the schema’s fields can serve as practical levers for screening, dataset growth, evaluation, and preference-driven MT for HC medical communication.

## 7 Conclusion and Future Work

We introduced an annotation framework for Haitian Creole medical translation that foregrounds audience-appropriate language varieties, overall acceptability, and clinical risk. Designed for compact expert teams, the schema emphasizes high inter-annotator agreement through binary decisions and conditional tags while remaining expressive enough to capture the sociolinguistic and clinical dimensions that matter in practice. Across three rounds of annotation and adjudication, we produced a gold set with fully resolved labels and a normalization pipeline that integrates directly into evaluation and modeling.

Using the adjudicated pool of 152 EN→HC pairs, a transparent classifier–labeller baseline showed that acceptability detection is already reliable, severity scores function as a practical triage signal with interpretable recall–effort trade-offs, and error typing is tractable once quality issues are flagged. By contrast, foreign-influence judgments and rare error types remain limited chiefly by data coverage rather than by schema design, underscoring where expansion is most needed.

These results carry methodological and applied implications. The schema separates acceptability from clinical risk and encodes audience design explicitly, avoiding the conflation of sociolinguistic variation with generic “fluency.” The pilot further demonstrates that calibrated probabilities and threshold-free metrics can support workflows that require prioritization and ranking rather than binary calls. Because error tags align with MQM, the framework bridges annotation, evaluation, and downstream modeling in a reproducible way.

Future work will expand and rebalance the dataset through targeted sampling across Haiti and diaspora communities, reinforce lexicographic grounding for contact-sensitive forms, and add document-level cues where linguistic variation and meaning unfold across sentences. The same fields are reward-ready: they can be used to train pref-

erence models and RLHF systems that steer MT outputs toward plain-language, culturally appropriate HC while down-weighting risky or foreign-influence renderings.

## References

- Hilal Al Shamsi, Abdullah G Almutairi, Sulaiman Al Mashrafi, and Talib Al Kalbani. 2020. Implications of language barriers for healthcare: a systematic review. *Oman medical journal*, 35(2):e122.
- Ron Artstein and Massimo Poesio. 2008. Inter-coder agreement for computational linguistics. *Computational linguistics*, 34(4):555–596.
- Derek Bickerton. 1973. The nature of a creole continuum. *Language*, pages 640–669.
- David Campbell. 2012. Cultural competency in haitian-serving community health centers in south florida. *McGill Journal of Medicine*, 14(1).
- Centers for Disease Control and Prevention. 2016. [Everyday words for public health communication](#). Report / plain language guidance, Centers for Disease Control and Prevention, Office of the Associate Director for Communication. Accessed: May 12, 2026.
- Michel DeGraff. 2005. Linguists’ most dangerous myth: The fallacy of creole exceptionalism. *Language in society*, 34(4):533–591.
- Yves Dejean. 2010. Creole and education in haiti. *The Haitian Creole language: History, structure, use, and education*, pages 199–216.
- Rohit Dholakia and Anoop Sarkar. 2014. Pivot-based triangulation for low-resource languages. In *Proceedings of the 11th Conference of the Association for Machine Translation in the Americas: MT Researchers Track*, pages 315–328.
- Chandrika Divi, Richard G Koss, Stephen P Schmalz, and Jerod M Loeb. 2007. Language proficiency and adverse events in us hospitals: a pilot study. *International journal for quality in health care*, 19(2):60–67.
- Hanze Dong, Wei Xiong, Bo Pang, Haoxiang Wang, Han Zhao, Yingbo Zhou, Nan Jiang, Doyen Sahoo, Caiming Xiong, and Tong Zhang. 2024. Rlhf workflow: From reward modeling to online rlhf. *arXiv preprint arXiv:2405.07863*.
- Dominique Fattier-Thomas. 1984. De la variété rèk à la variété swa: Pratiques vivantes de la langue en haïti. *Conjonction*161, 162:39–51.
- Glenn Flores, Milagros Abreu, Cara Pizzo Barone, Richard Bachur, and Hua Lin. 2012. Errors of medical interpretation and their potential clinical consequences: a comparison of professional versus ad hoc versus no interpreters. *Annals of emergency medicine*, 60(5):545–553.

- Glenn Flores, M Barton Laws, Sandra J Mayo, Barry Zuckerman, Milagros Abreu, Leonardo Medina, and Eric J Hardt. 2003. Errors in medical interpretation and their potential clinical consequences in pediatric encounters. *Pediatrics*, 111(1):6–14.
- Bryant C Freeman. 1997. *Haitian-English English-Haitian Medical Dictionary*. [Lawrence, Kan.]: Institute of Haitian Studies, University of Kansas.
- Sonia Halimi and Pierrette Bouillon. 2019. Google translate and babeldr in community medical settings: Challenges of translating into arabic. In *Arabic translation across discourses*, pages 27–44. Routledge.
- Benjamin Hebblethwaite. 2012. French and underdevelopment, haitian creole and development: Educational language policy problems and solutions in haiti. *Journal of Pidgin and Creole languages*, 27(2):255–302.
- Maude Heuretélou, Féquière Vilsaint, Erst Mirville, Michel-Ange Hyppolite, and John D. Nickrosz. 2000. *English / Haitian Creole Medical Dictionary*. Educa Vision.
- Patricia Hudelson and François Chappuis. 2024. Using voice-to-voice machine translation to overcome language barriers in clinical communication: an exploratory study. *Journal of General Internal Medicine*, 39(7):1095–1102.
- Kerry Hwang, Sue Williams, Emiliano Zucchi, Terence WH Chong, Monita Mascitti-Meuter, Dina LoGiudice, Anita MY Goh, Anita Panayiotou, and Frances Batchelor. 2022. Testing the use of translation apps to overcome everyday healthcare communication in australian aged-care hospital wards—an exploratory study. *Nursing open*, 9(1):578–585.
- Elizabeth A Jacobs, Donald S Shepard, Jose A Suaya, and Esta-Lee Stone. 2004. Overcoming language barriers in health care: costs and benefits of interpreter services. *American journal of public health*, 94(5):866–869.
- Leah S Karliner, Elizabeth A Jacobs, Alice Hm Chen, and Sunita Mutha. 2007. Do professional interpreters improve clinical care for patients with limited english proficiency? a systematic review of the literature. *Health services research*, 42(2):727–754.
- Erin Kobetz, Janelle Menard, Betsy Barton, Jennifer Cudris Maldonado, Joshua Diem, Pascale Denize Auguste, and Larry Pierre. 2010. Barriers to breast cancer screening among haitian immigrant women in little haiti, miami. *Journal of immigrant and minority health*, 12(4):520–526.
- Erin Kobetz, Janelle Menard, Joshua Diem, Betsy Barton, Jenny Blanco, Larry Pierre, Pascale D Auguste, Marie Etienne, and Cheryl Brewster. 2009. Community-based participatory research in little haiti: challenges and lessons learned. *Progress in Community Health Partnerships: Research, Education, and Action*, 3(2):133–137.
- Annika Kreienbrinck, Saskia Hanft-Robert, Alina Ioana Forray, Asithandile Nozewu, and Mike Mösko. 2025. Usability of technological tools to overcome language barriers in healthcare—a scoping review. *Archives of Public Health*, 83(1):52.
- Julia Kreutzer, Shahram Khadivi, Evgeny Matusov, and Stefan Riezler. 2018. Can neural machine translation be improved with user feedback? *arXiv preprint arXiv:1804.05958*.
- William Lewis. 2010. Haitian creole: How to build and ship an mt engine from scratch in 4 days, 17 hours, & 30 minutes. In *Proceedings of the 14th Annual conference of the European Association for Machine Translation*.
- Arle Lommel, Hans Uszkoreit, and Aljoscha Burchardt. 2014. Multidimensional quality metrics (mqm): A framework for declaring and describing translation quality metrics. *Tradumàtica*, (12):0455–463.
- Chenyang Lyu, Zefeng Du, Jitao Xu, Yitao Duan, Minghao Wu, Teresa Lynn, Alham Fikri Aji, Derek F Wong, Siyou Liu, and Longyue Wang. 2023. A paradigm shift: The future of machine translation lies with large language models. *arXiv preprint arXiv:2305.01181*.
- Janelle Menard, Erin Kobetz, Joshua Diem, Martine Lifleur, Jenny Blanco, and Betsy Barton. 2010. The sociocultural context of gynecological health among haitian immigrant women in florida: applying ethnographic methods to public health inquiry. *Ethnicity & Health*, 15(3):253–267.
- Ludovic Mompelat. 2025. Recommendations for overcoming linguistic barriers in healthcare: Challenges and innovations in nlp for haitian creole. In *Proceedings of the Third Workshop on Resources and Representations for Under-Resourced Languages and Domains (RESOURCEFUL-2025)*, pages 20–31.
- National Coordinating Council for Medication Error Reporting and Prevention (NCC MERP). 2022. *Ncc merp index for categorizing medication errors (2022 revision)*. Technical report, NCC MERP. Revised categorization of medication errors by severity.
- National Institutes of Health. 2025. Nih clear communication initiative. <https://www.nih.gov/institutes-nih/nih-office-director/office-communications-public-liaison/clear-communication>. Accessed: May 12, 2026.
- Eva Maria Noack, Jennifer Schulze, and Frank Müller. 2021. Designing an app to overcome language barriers in the delivery of emergency medical services: participatory development process. *JMIR mHealth and uHealth*, 9(4):e21586.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, and 1 others. 2022. Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35:27730–27744.

- Anita Panayiotou, Anastasia Gardner, Sue Williams, Emiliano Zucchi, Monita Mascitti-Meuter, Anita MY Goh, Emily You, Terence WH Chong, Dina Logiudice, Xiaoping Lin, and 1 others. 2019. Language translation apps in health care settings: Expert opinion. *JMIR mHealth and uHealth*, 7(4):e11316.
- Rafael Pérez-Escamilla, Jonathan Garcia, and David Song. 2010. Health care access among hispanic immigrants: ¿alguien está escuchando?[is anybody listening?]. *NAPA bulletin*, 34(1):47–67.
- Ellen R Ryan, Wesley E Hawkins, Marilyn Parker, and Michele J Hawkins. 2004. Perceptions of access to us health care of haitian immigrants in south florida. *Florida Public Health Review*, 1:30–35.
- Hervé Spechbach, Johanna Gerlach, Sanae Mazouri Karker, Nikos Tsourakis, Christophe Combescure, Pierrette Bouillon, and 1 others. 2019. A speech-enabled fixed-phrase translator for emergency settings: Crossover study. *JMIR medical informatics*, 7(2):e13167.
- David Tezil. 2022. On the influence of kreyòl swa: Evidence from the nasalization of the haitian creole determiner/la/in non-nasal environments. *Journal of Pidgin and Creole Languages*, 37(2):291–320.
- David Tézil. 2024. Sociolinguistic challenges and new perspectives on determining french speakers in creole communities: the case of haiti. *International Journal of the Sociology of Language*, 2024(288):177–207.
- Translators without Borders. 2020. Twb glossary for covid-19 (multilingual plain-language glossary). <https://translatorswithoutborders.org/covid-19/>. Accessed: May 12, 2026.
- U.S. Department of Health and Human Services. 2013. National standards for culturally and linguistically appropriate services in health and health care: A blueprint for advancing and sustaining clas policy and practice. Technical report, U.S. Department of Health and Human Services, Office of Minority Health, Washington, DC.
- Albert Valdman. 2015. *Haitian Creole: structure, variation, status, origin*. Equinox Publishing Limited.
- Albert Valdman, Marvin D Moody, and Thomas E Davies. 2017. *English-Haitian Creole Bilingual Dictionary*. iUniverse.
- Haoxiang Wang, Wei Xiong, Tengyang Xie, Han Zhao, and Tong Zhang. 2024. Interpretable preferences via multi-objective reward modeling and mixture-of-experts. *arXiv preprint arXiv:2406.12845*.
- World Health Organization. 2021. Use plain language: Communicating health information clearly. <https://www.who.int/about/communications/understandable/plain-language>. Accessed: May 12, 2026.

# Parser agreement and disagreement in L2 Korean UD: Implications for human-in-the-loop annotation

Hakyung Sung<sup>1</sup> Gyu-Ho Shin<sup>2</sup>

<sup>1</sup>Psychology, Rochester Institute of Technology

<sup>2</sup>Linguistics, University of Illinois Chicago

hksgla@rit.edu ghshin@uic.edu

## Abstract

We propose a simplified human-in-the-loop workflow for second language (L2) Korean morphosyntactic annotation by leveraging agreement between two domain-adapted parsers. We first evaluate whether parser agreement can serve as a proxy for annotation correctness by comparing it with independent human judgments. The results show strong correspondence between parser and human judgments, supporting the feasibility of semi-automatic L2-Korean UD annotation. Further analysis demonstrates that parser disagreements cluster in linguistically predictable domains such as grammatical-relation distinctions and clause-boundary ambiguity. While many disagreement cases are tractable for iterative model refinement, others reflect deeper representational challenges inherent in parsing and tagging L2-Korean corpora.

## 1 Introduction

Second language (L2) learner corpora consist of language samples produced by individuals acquiring an L2. Natural language processing (NLP) provides computational methods for extracting linguistic features (e.g., part-of-speech tags, grammatical relations). Together, these fields share the goal of empirically modeling L2 production at scale (Meurers, 2015). However, applying NLP tools to L2 corpora has traditionally been considered problematic for two reasons. First, general-purpose models trained on well-edited native/first-language corpora are assumed to perform poorly on non-canonical learner language (Plank, 2016). Second, this assumption has been difficult to verify due to the lack of learner-language benchmarks.

In response, the Universal Dependencies (UD) framework has emerged as a principled foundation for annotating morphosyntactic features in L2 corpora (Masciolini et al., 2025; Zeman, 2025). UD offers cross-linguistic comparability while main-

taining relative simplicity in morphosyntactic annotation and parsing (de Marneffe et al., 2021).

As UD-annotated learner corpora have become available, subsequent work has evaluated morphosyntactic models on such data, yielding mixed results. In high-resource languages such as L2 English, UD-based transformer models achieve over 90% F1 in part-of-speech (POS) tagging and dependency parsing (Kyle and Eguchi, 2024). In relatively under-resourced languages such as L2 Korean, performance varies by layer: morpheme-level tagging reaches F1 88%, while dependency parsing remains substantially lower (LAS 57%) (Sung and Shin, 2025b). Performance also varies by genre and proficiency. For instance, written learner data yield lower parsing accuracy than spoken data in L2 English (Kyle and Eguchi, 2024). In L2 Korean, higher-proficiency learners tend to produce longer and more syntactically complex sentences, which increase parsing difficulty; accordingly, proficiency is negatively correlated with dependency head accuracy ( $r = -0.26$ ) (Sung and Shin, 2025b). Importantly, both studies report improved performance after fine-tuning on L2-annotated data, suggesting that domain adaptation partially mitigates these limitations.

Taken together, for researchers seeking to leverage NLP-based annotations in L2 research, the question is increasingly no longer whether morphosyntactic models can be applied to learner corpora, but how they can be effectively integrated into annotation workflows. Because performance varies across languages, proficiency levels, and genres, fully automatic annotation may remain insufficient. However, automatic pre-annotation paired with human verification offers a practical alternative. In this context, the present study explores a simplified human-in-the-loop workflow for L2-Korean UD annotation, testing whether parser agreement can guide selective human review without compromising annotation reliability.

## 2 Related work

### 2.1 Use of morphosyntactic annotation in L2 learner corpora research

**Part-of-Speech tags:** Part-of-Speech (POS) tags have been widely employed in learner corpora research. In L2 English, for example, phraseological competence (i.e., the use of semi-/prefabricated expressions) has been examined through the automatic extraction of patterns from POS-tagged corpora (e.g., [Granger and Bestgen, 2014](#)). POS tags have also been used to measure lexical richness and disambiguate homographs in L2 Spanish ([Díez-Ortega and Kyle, 2024](#)). In the case of L2 Korean, language-specific POS tagsets enabled the representation of fine-grained morphemic distinctions within words ([Sung et al., 2024](#)).

**Dependency relations:** Syntactic information derived from dependency relations has likewise supported diverse corpus-based investigations. For instance, dependency-based phraseological units have been analyzed in L2 Dutch to examine the lexis–grammar interface ([Rubin et al., 2025](#)), while dependency representations have been used to assess lexical and syntactic complexity in L2 Russian ([Kisselev et al., 2022](#)). In L2 English, prior work examined n-grams within specific dependency relations (e.g., [Paquot, 2019](#)), verb–argument constructions and related predicate–argument patterns (e.g., [Kyle and Crossley, 2017](#)), and broader measures of syntactic complexity (e.g., [Kyle and Crossley, 2018](#)). Similarly, [Hao et al. \(2024\)](#) employed dependency parsing to investigate syntactic complexity in L2-Chinese writing.

### 2.2 Reliability of morphosyntactic annotation on L2 corpora

Although previous studies (as exemplified in Section 2.1) have reported important empirical findings based on extracted morphosyntactic features, their validity depends in part on annotation reliability. If automatic analyses are inaccurate, resulting conclusions may be compromised. While several studies evaluated the performance of NLP models on L2 corpora (e.g., [Berzak et al., 2016](#)), findings have been mixed and often limited in scope. As noted by [Kyle and Eguchi \(2024\)](#), many investigations have focused on isolated components (e.g., selected POS tags or dependency relations) rather than overall morphosyntactic performance. Moreover, earlier studies relied on neural architectures trained pri-

marily on well-edited standard-language data (e.g., news articles), without adaptation to L2 learner language.

A recent advance has been the adoption of domain adaptation techniques, in which annotated L2 treebanks are incorporated into model training to improve annotation quality ([Kyle and Eguchi, 2024](#); [Sung and Shin, 2025b](#)). Although effective, this approach presupposes the availability of reliable L2 annotations. This aspect makes it important to examine how such annotations are produced in existing L2 corpora.

### 2.3 UD annotation practices in L2 corpora

Over the past decade, an increasing number of L2 corpora have been annotated within the UD framework for diverse research purposes ([Zeman, 2025](#)). Our review identified eight such corpora to date. [Masciolini et al. \(2025\)](#) provide a detailed comparison, outlining their design characteristics (e.g., modality, size, annotation status) and the strategies adopted to address L2-specific phenomena, including ill-formed or non-canonical constructions.

Here, we examine the annotation methodologies underlying these corpora, focusing on whether morphosyntactic annotation was conducted either fully manually or through semi-automatic procedures (i.e., automatically annotated and then corrected by humans; see Table 1). Most UD-based L2 corpora relied on fully manual annotation, with relatively few adopting automatic approaches supplemented by human correction. While manual annotation supports quality control (e.g., through inter-annotator reliability), it is resource-intensive, difficult to scale, and challenging to replicate consistently across projects and annotator teams. Hybrid approaches that combine automatic processing with human oversight may therefore provide a more efficient and reproducible alternative.

### 2.4 Human-in-the-Loop annotation via model agreement

Human-in-the-Loop (HITL) machine learning broadly refers to workflows in which human expertise is intentionally integrated into automated systems to guide, validate, or correct model behavior ([Mosqueira-Rey et al., 2023](#)). Rather than replacing automation, such approaches strategically combine machine efficiency with human judgment, enhancing scalability while preserving reliability. They have been increasingly adopted in domains where full automation is unreliable or where high-

Language (domain)	Annotation method(s)	Reference
Chinese (written)	Manual	Lee et al. (2017)
English (written)	Manual	Berzak et al. (2016)
English (spoken)	Manual	Kyle et al. (2022)
Italian (written)	Semi-automatic	Di Nuovo et al. (2019)
Korean (written)	Manual; Semi-automatic	Sung and Shin (2024); Sung et al. (2025)
Russian (written)	Manual (single annotator)	Rozovskaya (2024)
Spanish (written)	Manual	Pulido et al. (2025)
Swedish (written)	Semi-automatic	Volodina et al. (2025)

Table 1: Overview of UD annotation practices in L2 learner corpora

stakes decisions require human oversight (Amershi et al., 2014; Holzinger, 2016).

Within the broader HITL taxonomy (Holmberg et al., 2020), active learning represents a prominent paradigm in which models select informative or uncertain instances for human annotation (Settles, 2009). Here, annotators function as oracles, and their feedback is used to iteratively refine the model. Active learning has been widely applied in NLP tasks such as POS tagging, dependency parsing, and text classification to reduce annotation cost while maintaining performance (Ringger et al., 2008).

Beyond uncertainty-based selection, prior work in dependency parsing has shown that simple ensemble strategies (i.e., agreement-based voting across multiple parsers) can produce robust predictions without requiring complex meta-modeling (Surdeanu and Manning, 2010). This suggests that model agreement can serve as a complementary signal of confidence. Building on this idea, the present study leverages parser agreement to guide selective human intervention, using disagreement cases as candidates for targeted review.

### 3 Experiment

In this exploratory study, we examine a simplified HITL workflow for L2-Korean UD annotation. Specifically, we compare the outputs of two independently fine-tuned parsers, treating agreement as a proxy for reliable annotation and disagreement as a signal for targeted human review. We assess whether such a setup can support more efficient annotation in future L2-Korean corpora. The study addresses the following research questions:

1. Can parser agreement reliably serve as a proxy for human annotation agreement?

Metric	Stanza	Trankit
LEMMA	95.64	88.84
XPOS	89.72	91.81
UAS	85.53	92.28
LAS	80.36	89.13

Table 2: Performance comparison (F1 scores) of fine-tuned Stanza and Trankit models on the L2K-UD test dataset

2. How much manual correction is required to resolve parser disagreements?
3. Which morphosyntactic categories exhibit disagreement, and how can these patterns inform annotation refinement?

#### 3.1 Proposed framework

The proposed annotation framework consists of three steps.

**Step 1: Automatic annotation.** Two domain-adapted parsers—*Stanza* (Qi et al., 2020) and *Trankit* (Van Nguyen et al., 2021)—were applied.<sup>1</sup> We examined four layers: LEMMA, XPOS, HEAD, and DEPREL.<sup>2</sup> Table 2 reports in-domain performance on the UD-KSL test set (Sung and Shin, 2025a).

**Step 2: Cross-model comparison.** Parser outputs were compared at the token level. Tokens with identical outputs were provisionally accepted, whereas any disagreement triggered human review.

<sup>1</sup>Stanza is a neural pipeline that performs joint tokenization, POS tagging, lemmatization, and dependency parsing using BiLSTM-based and transition-based components, while Trankit is a transformer-based multilingual pipeline built on XLM-R representations for joint morphosyntactic analysis. Both models were fine-tuned on the UD-KSL training set (Sung and Shin, 2025a), a learner corpus of L2 Korean writing annotated with morpheme-level segmentation, XPOS tags, and dependencies.

<sup>2</sup>UPOS was excluded because it is deterministically derived from XPOS under the current annotation scheme (Sung et al., 2025).

**Step 3: Human adjudication.** Two trained annotators independently reviewed the flagged tokens. If they assigned identical annotations at the token level, their decision was adopted as the gold label. In cases of disagreement, a third annotator (one of the authors) adjudicated by reviewing both model outputs and the independent annotations to assign the final label.

### 3.2 Evaluation of parser agreement

Given the exploratory nature of this study, we first collected fully independent annotations from both human annotators for all sentences before restricting review to parser-disagreement cases. This design allowed us to assess whether parser agreement could serve as a proxy for human annotation agreement. Specifically, we measured how often parser agreement coincided with inter-annotator agreement. Under the assumption that alignment is sufficiently strong (i.e., human agreement exceeds 90% within parser-agreement cases), parser-agreement cases could be retained automatically in future annotation rounds, enabling human effort to focus primarily on disagreement cases.

### 3.3 Pipeline validation

Prior to full-scale annotation of the target corpus, we conducted a small-scale validation experiment to evaluate whether the proposed semi-automatic pipeline introduced annotation noise. We randomly sampled 500 L2-Korean sentences from the KoLLA corpus (Lee et al., 2009), none of which had been used to train the fine-tuned models. Annotation proceeded incrementally. In each round, a batch of 100 sentences was annotated using the framework described in Section 3.1, with human review limited to tokens where the two models disagreed. Adjudicated annotations were then incorporated into the training data, and both models were fine-tuned for 10 epochs on the expanded dataset. After each round, performance was evaluated on the UD-KSL test set to determine whether accuracy improved, stabilized, or declined. This procedure was repeated for five rounds.

Figure 1 presents model performance on the test set across incremental fine-tuning rounds. Overall, performance remained stable: both Stanza and Trankit maintained accuracies of approximately 85% or higher, with no observable decline. These results indicate that the proposed pipeline did not introduce substantial degradation in annotation quality during incremental fine-tuning.

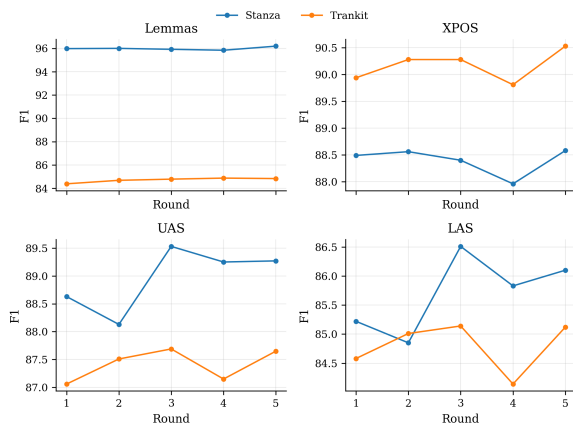


Figure 1: Model performance on the test set across fine-tuning rounds

## 4 Results

### 4.1 Annotations

Two trained annotators annotated a total of 2,208 sentences drawn from argumentative essays written by Japanese and English learners of Korean over a one-month period. We adopted the morphosyntactic annotation scheme developed in prior L2 Korean UD annotation work, most recently detailed in Sung et al. (2025).

### 4.2 Parser agreement as a proxy for human annotation agreement

To evaluate whether parser agreement can serve as a proxy for human annotation agreement, we calculated token-level agreement rates, excluding punctuation. Table 3 summarizes the results. Across all features, the two parsers agreed on 82% of token-level decisions. Within these parser-agreement cases, the human annotators also agreed on 93% of instances, with agreement exceeding 90% for all four features. These findings indicate that parser agreement closely corresponded to independent human agreement, supporting its use as a practical proxy for human annotation agreement in the HITL workflow.

Feature	Model agreement	Human agreement
LEMMA	78.09	96.97
XPOS	84.83	91.97
HEAD	78.62	90.22
DEPREL	88.78	92.55
Average	82.58	92.93

Table 3: Token-level agreement between the two parsers and corresponding human annotator agreement rates (punctuation excluded)

### 4.3 Human intervention following model mismatch

Based on the agreement results, we next evaluated a workflow in which human annotators intervened only when the two models disagreed.<sup>3</sup> At the token level, 7,994 out of 25,814 tokens (31%) required human correction in at least one morphosyntactic feature.<sup>4</sup>

Feature-level correction counts are provided in Table 4, and feature-level adjudication corrections are summarized in Table 5. A total of 2,019 out of 25,814 tokens (8%) required further modification after initial review. These cases reflect instances in which the two annotators did not converge, necessitating adjudication by a third annotator.

Feature	Human	Total	Rate (%)
LEMMA	4,485	25,814	17.37
XPOS	2,263	25,814	8.77
HEAD	3,798	25,814	14.71
DEPREL	1,713	25,814	6.64

Table 4: Feature-level human corrections following model mismatches

Feature	Fixed	Total	Rate (%)
LEMMA	1,125	25,814	4.36
XPOS	1,461	25,814	5.66
HEAD	581	25,814	2.25
DEPREL	1,153	25,814	4.47

Table 5: Feature-level third-annotator corrections

Overall, these findings demonstrate how parser disagreement can structure a tiered annotation workflow. Restricting human review to model-disagreement cases substantially reduces effort, with nearly 70% of tokens requiring no intervention after alignment. These tokens likely represent morphosyntactic categories that are relatively stable and well captured by the fine-tuned models. Most remaining disagreement cases were resolved through agreement between two annotators, suggesting that they are tractable.

In contrast, the 8% of tokens requiring third-annotator adjudication represent persistent disagreement across both models and trained annotators. These instances often involved structurally

<sup>3</sup>Prior to analysis, tokenization mismatches were resolved to ensure proper token-level alignment. All reported agreement and intervention rates are based on the aligned data.

<sup>4</sup>Because a single token may require correction in multiple features, counts are not mutually exclusive.

complex or potentially ambiguous linguistic units. Although tentative, such residual disagreement at high overall accuracy may reflect not only model limitations but also indeterminacy in linguistic categories or variation in annotation conventions (Manning, 2011).

### 4.4 Disagreement analysis

To further characterize these disagreement patterns, we analyzed where and why the models diverged.

#### 4.4.1 Distribution of dependency-relation disagreements

We first conducted a focused analysis of disagreement patterns in the dependency-relation (DEPREL) layer, as dependency-relation labeling exhibited relatively high disagreement rates. Disagreements were classified according to the primary syntactic decision involved: (1) grammatical relation identification, (2) clause-boundary and clause-type differentiation, (3) discourse-level structural organization, and (4) modifier attachment.

- **Grammatical relation identification:** This category captures instability in assigning core grammatical relations (e.g., subject, object, oblique), including contrasts such as `nsubj-obj`, `nsubj-obl`, and `obj-obl`.
- **Clause boundary:** This category reflects uncertainty in clause typing and hierarchical embedding, including distinctions among adjectival (relative), adverbial, and complement clauses (e.g., `acl-advcl`, `advcl-ccomp`, and `advcl-root`).
- **Discourse-level organization:** This category involves higher-level decisions at the syntax-discourse interface, such as coordination scope, root status, and left dislocation (e.g., topic-marked elements). Recurrent contrasts included `dislocated-nsubj`, `root-conj`, and `conj-advcl`.
- **Modifier attachment:** This category captures ambiguity in determining the structural status or attachment site of modifiers, with contrasts such as `amod-acl` and `nmod-obl`.

Table 6 summarizes mismatch frequencies across categories. To illustrate these patterns more concretely, representative examples from the annotated texts are provided below.<sup>5</sup>

<sup>5</sup>Sentences have been streamlined for clarity.

Mismatch type	Count
Grammatical relation	263
Clause boundary	206
Discourse / structure	235
Modifier attachment	152

Table 6: Major categories of dependency-relation disagreements between the two parsers. Counts indicate the total number of mismatches in each category.

Grammatical-relation ambiguity (e.g., nsubj-obj, nsubj-obl) accounts for the largest share of mismatches. These alternations typically arise when case-marking is underspecified or omitted, obscuring whether a nominal is analyzed as subject, object, or oblique. For example, when nominative (Figure 2) or accusative case markers (Figure 3) are dropped, a preverbal noun phrase may be ambiguously interpreted, leading to divergent grammatical relation identifications across parsers.

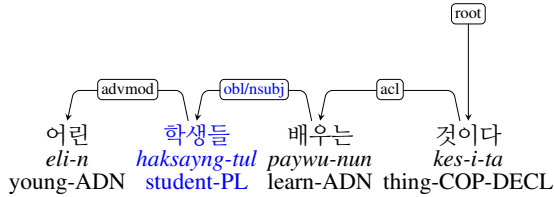


Figure 2: Grammatical-relation ambiguity under case marker omission 1. The nominal 학생들 ‘students’ was tagged as either obl or nsubj across the parsers; contextual interpretation favors nsubj. (Translated as ‘(It) is that young students learn.’)

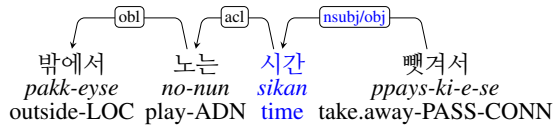


Figure 3: Grammatical-relation ambiguity under case marker omission 2. The nominal 시간 ‘time’ received conflicting tags (nsubj vs. obj); contextual interpretation favors obj. (Translated as ‘(When) time spent playing outside is taken away.’)

Clause-boundary ambiguity (e.g., acl-advcl, advcl-ccomp) reflects uncertainty in clause typing and hierarchical embedding. Figure 4 illustrates a case in which a clause can be analyzed either as an adnominal modifier (top) or as a subordinate clause within the predicate domain (bottom).

Compared to grammatical-relation ambiguity, clause-boundary ambiguity poses greater challenges for two reasons. First, the polyfunctionality

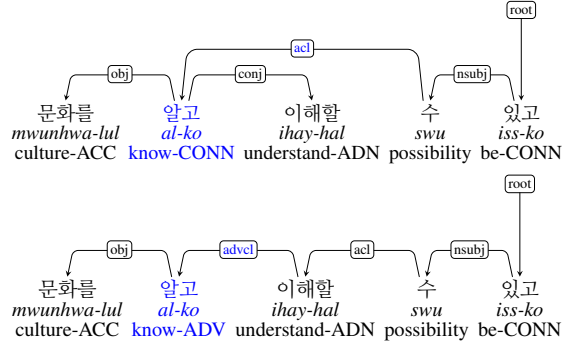


Figure 4: Clause-boundary ambiguity 1. The connective verb 알고 (‘know-CONN/ADV’) is ambiguous. In the first analysis, the clause is treated as an adnominal modifier (acl); in the second, it is analyzed as an adverbial clause (advcl); the appropriate annotation cannot be determined from the sentence in isolation. (Translated as Top: ‘[One] can know and understand the culture.’ Bottom: ‘After knowing the culture, [one] can understand it.’)

of the connective 고 -ko in Korean frequently triggers disagreement, as its interpretation depends on discourse-semantic cues rather than overt syntactic marking. Second, clause-typing uncertainty often interacts with higher-level structural mismatches (e.g., conj-advcl), corresponding to the third category of disagreement. Such cases therefore require adjudication informed by broader sentential or discourse context.

Meanwhile, not all clause-type disagreements arise from context-dependent ambiguity. In advcl-ccomp mismatches, some cases instead appear to reflect model difficulty in learning complement structures headed by the -지 (-ci) complementizer (Figure 5).

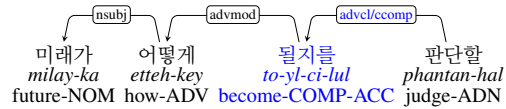


Figure 5: Clause-type disagreement 2. The embedded clause 될지를 (‘how [it] will become’) functions as a clausal complement (ccomp) of the matrix predicate 판단할 (‘judge’), a pattern that one parser consistently failed to capture; morphosyntactic cues (i.e., the complementizer 지 and accusative marker 를) favor a ccomp. (Translated as: ‘[One] can judge how the future will turn out.’)

Discourse-related ambiguity (e.g., dislocated-nsubj) reflects instability at the syntax-discourse interface, particularly in topic-prominent constructions where left-dislocated elements may be misanalyzed as canonical subjects.<sup>6</sup>

<sup>6</sup>Examples of structural mismatches (e.g., root-conj)

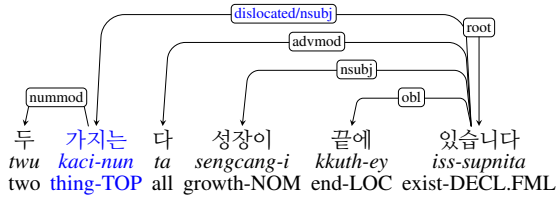


Figure 6: Discourse-level misanalysis. The topic-marked noun phrase 가지는 (‘things-TOP’) functions as dislocated, a pattern that one parser consistently failed to capture. (Translated as ‘As for the two things, both ultimately result in growth.’)

Finally, modifier attachment ambiguity (e.g., amod-ac1, nmod-ob1) reflects uncertainty in hierarchical scope, particularly when linear proximity does not clearly determine attachment. As illustrated in Figure 7, an adnominal form can be analyzed either as a lexical adjectival modifier (amod) or as a reduced relative clause (ac1). Similarly, Figure 8 shows that a locative phrase can attach either to a noun phrase (nmod) or to the predicate as a clausal oblique (obl), depending on its interpreted scope. Such modifier-attachment ambiguity is likewise context-dependent, requiring broader interpretive information beyond local morphosyntactic cues.

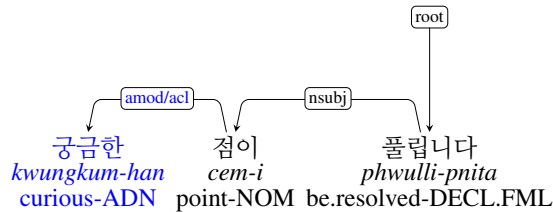


Figure 7: Modifier attachment ambiguity 1. The form 궁금한 (‘curious-ADN’) can function either as a lexical adjectival modifier (amod) or as a reduced relative clause (ac1) modifying cem (‘point’). (Translated as ‘The questions are resolved.’)

#### 4.4.2 Morphological-level disagreements

Table 7 presents the twenty most frequent morpheme-level XPOS mismatches. Similar to dependency relations, these disagreements form recurrent patterns rather than occurring randomly. First, many mismatches involved case particles (e.g., JKS, JKB, JKO, JKC, JX, JKG). These contrasts often reflect functional ambiguity, particularly in distinguishing structural case markers from auxiliary or semantic/discourse particles. Such morphological ambiguity closely parallels the dependency-level

are not presented separately, as they typically co-occur with clause-boundary ambiguities discussed above.

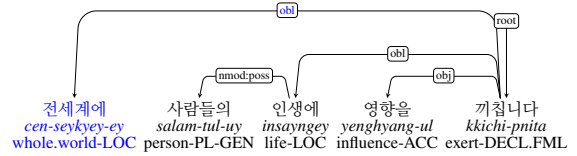
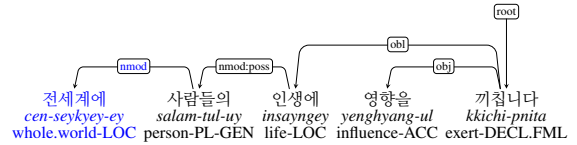


Figure 8: Modifier attachment ambiguity in a locative phrase. The locative 전세계에 (‘whole.world-LOC’) attaches either to 사람들 (‘people’) as nmod or to 끼칩니다 (‘exert’) as obl. Despite the locative marker (-에), the first analysis yields a possessive-like reading (‘people of the whole world’). (Translated as Top: ‘[It] affects the lives of people in the whole world.’ Bottom: ‘In the whole world, [it] affects people’s lives.’)

Rank	Stanza	Trankit	Count
1	NNG+JKS	NNG+JKC	54
2	NNG+JC	NNG+JKB	48
3	NNG+XSA+ETM	XR+XSA+ETM	40
4	VV+ETM	NV+ETM	28
5	NNG+VCP+EF	NNB+VCP+EF	27
6	VV+EC	NV+EC	25
7	-	JX	23
8	MAG	NNG	23
9	VV+ETM	VX+ETM	22
10	NNG+JKB	NF+JKB	22
11	VX+EC	VV+EC	21
12	NNG+XSA+EC	XR+XSA+EC	20
13	NF+JKO	NNG+JKO	20
14	NNG	MAG	19
15	NNG+NNG+NNG+JKG	NNG+NNG+NNG+JX	18
16	NNG+JKO	NF+JKO	18
17	VV+ETM+NNB	VV+ETM+NNB+JX	18
18	VV+EC+VX	VV+EC+VX+EC	17
19	VA+EC	VV+EC	17
20	VV+EC+VX	VV+EC	17

Table 7: Twenty most frequent XPOS disagreement pairs between Stanza and Trankit. Counts indicate the number of tokens assigned different morpheme-level POS analyses.

ambiguities observed in grammatical-relation identification and discourse structure.

Second, high-frequency mismatches often arose from differences in lexical decomposition and root identification (e.g., NNG vs. XR; NNG vs. NF; VV vs. NV). Two issues are implicated. First, the distinction between common nouns (NNG) and lexical roots (XR) is not always clear-cut in Korean, as root classification can be inherently ambiguous. Second, this ambiguity is further amplified in learner language, where non-canonical forms are frequent. In our annotation scheme, tags such as NF, NV, and NA mark ill-formed or irregular forms. While parsers may recover intended lexical items for recurring spelling

errors through dictionary-based matching, novel or idiosyncratic errors often lack lexical support, resulting in divergent analyses. These patterns underscore the need for more systematic approaches to learner-specific morphological variation.

Finally, some disagreements involved segmentation differences, including the insertion or omission of functional morphemes (e.g., additional JX or EC). These cases reflect variation in morphological parsing strategies rather than simple tagging errors.

## 5 Conclusion

The purpose of this study was to evaluate a simplified HITL workflow for L2-Korean UD annotation. The findings provide three main implications, which may be relevant for researchers working on morphosyntactic annotation in learner corpora.

First, across all annotation features, the two independently domain-adapted parsers agreed on 82% of token-level decisions. Within these consensus cases, human annotators also agreed on 93% of instances. This strong alignment suggests that parser agreement reliably predicts correspondence with independent human judgments. For scalable L2 annotation, parser consensus may therefore serve as an effective filtering mechanism, substantially reducing the need for exhaustive manual verification. In addition, as noted by one reviewer, this binary agreement approach could be extended within an ensemble framework (e.g., [Surdeanu and Manning, 2010](#)), which may enable more robust consensus estimation.

Second, despite high overall agreement, 31% of tokens required human review in at least one feature, and 8% required adjudication after initial correction. These findings suggest that morphosyntactic disagreement operates at multiple levels. Some cases are readily resolved through annotator agreement and are amenable to iterative model refinement, whereas others reflect deeper representational challenges in assigning L2-Korean forms to discrete morphosyntactic categories.

Third, parser disagreements clustered in linguistically predictable domains rather than occurring randomly. Our analysis showed that many involved argument-role distinctions and complement structures headed by complementizers, suggesting that targeted sampling and focused retraining could improve performance. In contrast, clause-boundary and modifier-attachment ambiguities were often context-dependent, indicating that some disagree-

ments cannot be resolved through local morphosyntactic cues alone and may require broader contextual modeling. At the morphological level, frequent mismatches involved root identification, learner-specific spelling variation, and segmentation differences. These patterns highlight the need for systematic strategies to handle learner-generated forms, particularly when such non-standard forms are not attested in the training data or lexical resources.

In conclusion, we examined whether parser agreement can serve as a principled triaging mechanism in L2 annotation. The results point to its potential, while also highlighting the multi-level nature of morphosyntactic disagreement. Distinguishing between tractable modeling limitations and deeper representational ambiguities remains important for achieving efficient yet reliable analysis of learner language.

## Limitations

First, the dataset consists exclusively of argumentative writing by adult L2-Korean learners. Because parser performance may vary by proficiency, age, genre, and language background, the generalizability of these findings is limited.

Second, although we briefly noted issues related to spelling-error tags, this study did not systematically examine learner-specific morphological variation. Developing principled approaches to modeling such variation is therefore an important direction for future research.

Third, the proposed workflow reflects a simplified HITL design rather than a fully integrated, interactive system. For instance, it did not incorporate dynamic confidence estimation, active learning, or real-time model updating. Research in this line would benefit from incorporating these aspects into the HITL design.

## Acknowledgments

This study was supported by the 2024 Korean Studies Grant Program of the Academy of Korean Studies (AKS-2024-R-012). The authors gratefully acknowledge Youkyung Sung and Chanyoung Lee for their contributions to manual annotation, and Jeong Eun Shin for providing the data.

## References

- Saleema Amershi, Maya Cakmak, William Bradley Knox, and Todd Kulesza. 2014. [Power to the people: The role of humans in interactive machine learning](#). *AI magazine*, 35(4):105–120.
- Yevgeni Berzak, Jessica Kenney, Carolyn Spadine, Jing Xian Wang, Lucia Lam, Keiko Sophie Mori, Sebastian Garza, and Boris Katz. 2016. [Universal Dependencies for learner English](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 737–746, Berlin, Germany. Association for Computational Linguistics.
- Marie-Catherine de Marneffe, Christopher D. Manning, Joakim Nivre, and Daniel Zeman. 2021. [Universal dependencies](#). *Computational Linguistics*, 47(2):255–308.
- Elisa Di Nuovo, Cristina Bosco, Alessandro Mazzei, and Manuela Sanguinetti. 2019. [Towards an italian learner treebank in universal dependencies](#). In *Proceedings of the Sixth Italian Conference on Computational Linguistics (CLiC-it 2019)*, pages 151–158, Bari, Italy. CEUR Workshop Proceedings.
- María Díez-Ortega and Kristopher Kyle. 2024. [Measuring the development of lexical richness of l2 spanish: A longitudinal learner corpus study](#). *Studies in Second Language Acquisition*, 46(1):169–199.
- Sylviane Granger and Yves Bestgen. 2014. [The use of collocations by intermediate vs. advanced non-native writers: A bigram-based study](#). *International Review of Applied Linguistics in Language Teaching*, 52(3):229–252.
- Yuxin Hao, Xuelin Wang, Shuai Bin, Qihao Yang, and Haitao Liu. 2024. [How syntactic complexity indices predict chinese l2 writing quality: An analysis of unified dependency syntactically-annotated corpus](#). *Assessing Writing*, 61:100847.
- Lars Holmberg, Paul Davidsson, and Per Linde. 2020. [A feature space focus in machine teaching](#). In *2020 IEEE International Conference on Pervasive Computing and Communications Workshops (PerCom Workshops)*, pages 1–2.
- Andreas Holzinger. 2016. [Interactive machine learning for health informatics: when do we need the human-in-the-loop?](#) *Brain informatics*, 3(2):119–131.
- Olesya Kisselev, Rossina Soyán, Dmitrii Pastushenkov, and Jason Merrill. 2022. [Measuring writing development and proficiency gains using indices of lexical and syntactic complexity: Evidence from longitudinal russian learner corpus data](#). *The Modern Language Journal*, 106(4):798–817.
- Kristopher Kyle and Scott Crossley. 2017. [Assessing syntactic sophistication in l2 writing: A usage-based approach](#). *Language Testing*, 34(4):513–535.
- Kristopher Kyle and Scott A Crossley. 2018. [Measuring syntactic complexity in l2 writing using fine-grained clausal and phrasal indices](#). *The Modern Language Journal*, 102(2):333–349.
- Kristopher Kyle and Masaki Eguchi. 2024. [Evaluating nlp models with written and spoken l2 samples](#). *Research Methods in Applied Linguistics*, 3(2):100120.
- Kristopher Kyle, Masaki Eguchi, Aaron Miller, and Theodore Sither. 2022. [A dependency treebank of spoken second language English](#). In *Proceedings of the 17th Workshop on Innovative Use of NLP for Building Educational Applications (BEA 2022)*, pages 39–45, Seattle, Washington. Association for Computational Linguistics.
- John Lee, Herman Leung, and Keying Li. 2017. [Towards universal dependencies for learner chinese](#). In *Proceedings of the NoDaLiDa 2017 Workshop on Universal Dependencies (UDW 2017)*, volume 135 of *Linköping Electronic Conference Proceedings*, pages 67–71. Linköping University Electronic Press.
- Sun-Hee Lee, Seok Bae Jang, and Sang-Kyu Seo. 2009. [Annotation of korean learner corpora for particle error detection](#). *Calico Journal*, 26(3):529–544.
- Christopher D Manning. 2011. [Part-of-speech tagging from 97% to 100%: is it time for some linguistics?](#) In *International conference on intelligent text processing and computational linguistics*, pages 171–189. Springer.
- Arianna Masciolini, Aleksandrs Berdičevskis, Maria Irena Szawerna, and Elena Volodina. 2025. [Annotating second language in universal dependencies: a review of current practices and directions for harmonized guidelines](#). In *Proceedings of the Eighth Workshop on Universal Dependencies (UDW, SyntaxFest 2025)*, pages 153–163.
- Detmar Meurers. 2015. [Learner corpora and natural language processing](#). In Sylviane Granger, Gaëtanelle Gilquin, and Fanny Meunier, editors, *The Cambridge Handbook of Learner Corpus Research*, pages 537–566. Cambridge University Press, Cambridge.
- Eduardo Mosqueira-Rey, Elena Hernández-Pereira, David Alonso-Ríos, José Bobes-Bascarán, and Ángel Fernández-Leal. 2023. [Human-in-the-loop machine learning: a state of the art](#). *Artificial Intelligence Review*, 56(4):3005–3054.
- Magali Paquot. 2019. [The phraseological dimension in interlanguage complexity research](#). *Second language research*, 35(1):121–145.
- Barbara Plank. 2016. [What to do about non-standard \(or non-canonical\) language in nlp](#). *arXiv:1608.07836 [cs]*.
- Emiliana Pulido, Robert Pugh, and Zoey Liu. 2025. [I speak for the árboles: Developing a dependency treebank for spanish l2 and heritage speakers](#). In

*Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 4: Student Research Workshop)*, pages 814–822.

Peng Qi, Yuhao Zhang, Yuhui Zhang, Jason Bolton, and Christopher D Manning. 2020. [Stanza: A python natural language processing toolkit for many human languages](#). *arXiv preprint arXiv:2003.07082*.

Eric Ringger, Marc Carmen, Robbie Haertel, Kevin Seppi, Deryle Lonsdale, Peter McClanahan, James Carroll, and Noel Ellison. 2008. [Assessing the costs of machine-assisted corpus annotation through a user study](#). In *Proceedings of the Sixth International Conference on Language Resources and Evaluation (LREC 2008)*, Marrakech, Morocco. European Language Resources Association (ELRA).

Alla Rozovskaya. 2024. [Universal Dependencies for learner Russian](#). In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 17112–17119, Torino, Italia. ELRA and ICCL.

Rachel Rubin, Bram Bulté, Magali Paquot, and Alex Housen. 2025. [Exploring complexity at the lexis-grammar interface: Diversity and sophistication of verb-argument structures in l2 dutch writing](#). *Journal of Second Language Writing*, 67:101183.

Burr Settles. 2009. Active learning literature survey. (1648). Computer Sciences Technical Report.

Hakyung Sung, Sooyeon Cho, and Kristopher Kyle. 2024. [An empirical evaluation of lexical diversity indices in l2 korean writing assessment](#). *Language Assessment Quarterly*, 21(2):159–180.

Hakyung Sung and Gyu-Ho Shin. 2024. [Constructing a dependency treebank for second language learners of korean](#). In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 3747–3758, Torino, Italia. ELRA and ICCL.

Hakyung Sung and Gyu-Ho Shin. 2025a. [Second language korean universal dependency treebank v1. 2: Focus on data augmentation and annotation scheme refinement](#). In *Proceedings of the Third Workshop on Resources and Representations for Under-Resourced Languages and Domains (RESOURCEFUL-2025)*, pages 13–19.

Hakyung Sung and Gyu-Ho Shin. 2025b. [Towards robust morphosyntactic analysis of L2 korean: Evaluating and fine-tuning a korean language model](#). *ACM Transactions on Asian and Low-Resource Language Information Processing*, 24(4):1–20.

Hakyung Sung, Gyu-Ho Shin, Chanyoung Lee, You Kyung Sung, and Boo Kyung Jung. 2025. [UD-KSL treebank v1.3: A semi-automated framework for aligning XPOS-extracted units with UPOS tags](#). In

*Proceedings of the 19th Linguistic Annotation Workshop (LAW-XIX-2025)*, pages 115–125, Vienna, Austria. Association for Computational Linguistics.

Mihai Surdeanu and Christopher D Manning. 2010. [Ensemble models for dependency parsing: cheap and good?](#) In *Human Language Technologies: The 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics*, pages 649–652.

Minh Van Nguyen, Viet Dac Lai, Amir Pouran Ben Veyseh, and Thien Huu Nguyen. 2021. [Trankit: A light-weight transformer-based toolkit for multilingual natural language processing](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: System Demonstrations*, pages 80–90.

Elena Volodina, Arianna Masciolini, Beáta Megyesi, Julia Prentice, Lisa Rudebeck, Gunlög Sundberg, and Mats Wirén. 2025. [SweLL with pride: How to put a learner corpus to good use](#). *Huminfra Handbook (forthcoming)*.

Daniel Zeman. 2025. [Corpus-based language comparison: From morphology to dependencies and beyond](#). *Estudos Linguísticos (São Paulo. 1978)*, 54(1):259–275.

## A Gloss abbreviations

Abbreviations used in interlinear glosses follow standard Leipzig conventions.

Abbreviation	Meaning
ACC	Accusative
ADN	Adnominal
ADV	Adverbial
CONN	Connective ending
COP	Copula
DECL	Declarative
FML	Formal speech level
LOC	Locative
NOM	Nominative
PASS	Passive
PL	Plural
TOP	Topic

# Rules-based system for Czech legal text readability

Kateřina Motalík Hodková and Ivan Kraus and Barbora Hladká

Institute of Formal and Applied Linguistics  
Faculty of Mathematics and Physics, Charles University  
Prague, Czech Republic  
{hodkova, hladka}@ufal.mff.cuni.cz, ivakra@centrum.cz

## Abstract

In this paper, we present a set of linguistic rules, employed to enhance the readability of legal texts. The rules were compiled and implemented as a rule-based module of PONK, an advisory tool that contributes to simplification and higher clarity of Czech legal texts, especially those intended for non-expert audience. Based on recurring phenomena in authentic texts and relevant scientific sources, the rules cover mainly the domains of syntax and lexicon. In addition, we present the results of application of the rules to a corpus of authentic legal texts, evaluated by a human annotator, and examine their impact.

## 1 Introduction

Legal texts can pose a significant challenge, particularly for readers without legal training. Their complexity arises from several factors, including long and complicated syntactic structures, legal terminology, and the use of archaic, ambiguous, or semantically vague expressions. Although such features can often be linked to historical or stylistic conventions of legal drafting, they can substantially hinder comprehension, and individuals affected by legal documents (e.g., citizens bound by laws of a country, parties of a contract, or parties of a court process) may struggle to interpret legal texts accurately, which may result in grave consequences. The readability and clarity of legal texts, therefore, presents an important objective, both from the perspective of legal certainty and from the general perspective of accessibility and transparency.

As the linguistic obstacles of legal texts are strongly language-dependent and potentially related to specific drafting conventions, we decided to focus on legal texts written in Czech. The paper presents both the design of the rule set (in the context of Czech grammar) and its application to authentic Czech legal texts, which is incorporated within the PONK tool. PONK tool is an advisory

tool designed to identify difficult-to-understand passages of legal texts (the tool does not automatically generate simplified versions of the texts). To assess the impact of the proposed rules, the identified segments are subjected to human annotation. Its purpose is to judge whether modifications of text segments indicated by PONK tool would contribute to an increased readability of the text or not. This annotation allows us to estimate the practical effectiveness of the rules and identify areas for further refinement.

In addition, our aim is to formulate the rules and the overall approach in a way that allows the methodology to be adapted to other languages with appropriate linguistic adjustments.

## 2 Legal text and related works

Legal language is a type of language for special purposes belonging to the standard register of a given language (Hodková et al., 2021). It is the language of legal text (written or spoken) and its objective is to transfer information belonging to the legal domain. Among the specialty languages, legal language occupies a particular position: whereas most specialty languages (be it in the domain of biology, astrophysics, or medicine) concern mostly only the given group of experts, legal language and legal texts involve virtually everyone via laws, contracts, etc. According to Baldinger (1984), specialty languages are, unlike natural languages, motivated and not arbitrary.

Terminological units constitute an inherent part of legal texts. Terms are lexical units designating legal concepts, abstract semantic units belonging to the domain of law (Cornu, 2005; Tomášek, 2003). Legal concepts often represent an obstacle in understanding legal text (mainly for readers without legal education). However, given that legal concepts are precisely defined (whether it is a legal definition or legal citation, Hodková et al. (2021)), they cannot

be altered, because this could obscure the intended meaning.

In anglophone countries, the tendencies to simplify the language of legal texts have a long tradition, concerning texts from jury instructions to contracts (Charrow and Charrow, 1979; Chromá, 2016; Diamond et al., 2012; Hartig and Lu, 2014; Martínez et al., 2022), including attempts to establish readability metrics for legal texts (Han et al., 2024). By contrast, in the Czech Republic, the readability of legal or administrative texts has been examined only recently (Šamánková and Kubíková, 2022; Bohuslav Halfar and Bučková, 2022; Chromý et al., 2021).

As Šamánková and Kubíková (2022) mention, a quality yet clear legal text should have a clear structure, should include unambiguous instructions for the addressee and all relevant information, should be brief, legally precise and correct, and written in simple and easy-to-understand language. Tomášek (2003) and Chromá (2016) list comparable properties. As studies show (Team, 2014), the vast majority (up to 80%) of addressees (general public or experts) prefer legal or administrative texts written in a simple and clear manner. Martínez et al. (2022) show that legal texts containing features such as legal jargon of low frequency and passive voice are more difficult to comprehend, while noting that some of the features pose greater difficulties than others. In addition, the British government suggests writing in such a way that would be suitable to “9 year old reading age”.<sup>1</sup>

Šamánková and Kubíková (2022) present a manual designed for authors of official administrative or legal texts, especially texts addressing the general public without extensive (or any) legal knowledge. It advises, among other things, to avoid an abundance of verbal nouns, passive voice, loanwords, and negations, and to eliminate, if possible, redundant expressions, repetitions, and abstract expressions. Similarly, Šváb (2023) advises to avoid synonyms or passive voice and to pay attention to the length of sentences, the distance between verbs and subjects, and the clear structure of the text.

In terms of regulations for the drafting of legislative texts, few sources can be found. For example, Důvodová zpráva k občanskému zákoníku (roughly translated as “Explanatory Memorandum to the Civil Code”) is relatively brief while commenting

<sup>1</sup><https://www.gov.uk/guidance/content-design/writing-for-gov-uk>

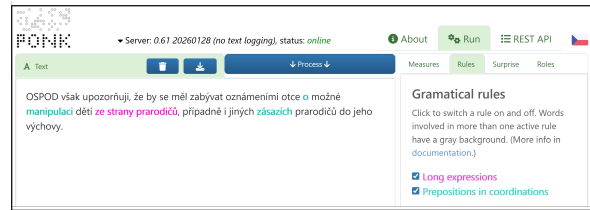


Figure 1: Front-end interface of the PONK’s rule-based module

on the Civil Code in this regard.<sup>2</sup> It mentions that one stipulation should not contain more than two paragraphs, while one paragraph should not consist of more than two sentences (but it does not address the length of sentences). Homonyms and polysemous words are supposed to be avoided, as well as multiple different terms for an identical concept (neither of which is strictly adhered to, see Hodková (2022)). The Memorandum also mentions that loanwords are not desirable while (vaguely defined) exceptions are allowed (see Hodková (2020) for a more detailed study on this topic). It also briefly defines the usage of some expressions (in the context of the law in question).

In practice, readable legal texts differ from less readable ones in linguistic features, although linguistic readability advice is not always easy to follow (Kraus, 2025). Despite a wide range of factors contributing to (un)readability, language remains an important component according to focus groups (Veřejný ochránce práv, 2024).

### 3 Overview of the PONK tool

PONK is a client-oriented tool that addresses the clarity and readability of legal texts (Novotná et al., 2025).<sup>3</sup> It is intended for authors of legal texts. The objective of PONK is to identify phenomena that hinder readability of the text and signal them to the author. In addition to the rules presented in this paper, PONK also incorporates additional modules for readability assessment. These include (1) a lexical surprise module, which analyzes the probability of occurrence of a given token in the given context (compared to other possible tokens); e.g., a high level of lexical surprise indicates a sudden change of context, and (2) a speech acts module, which

<sup>2</sup><http://obcanskyzakonik.justice.cz/images/pdf/Duvodova-zprava-NOZ-konsolidovana-verze.pdf>.

<sup>3</sup>The name itself is an acronym of “Psaní Orientované Na Klienta”, approximately translated as “Client-oriented writing”. The PONK tool is available at <https://quest.ms.mff.cuni.cz/ponk/>.

distinguishes between argumentative and normative texts and assigns to each of these two genres its typical inventory of so-called Speech Acts or Rhetorical Roles, that are based on the theory of legal writing. PONK also offers a metrics overview of the entire text, such as vocabulary, ratio and distance of verbs, or common readability metrics.<sup>4</sup> Figure 1 shows the front-end of the PONK’s rule-based module.

PONK rules are applied to texts processed with UDPipe 2, a tool that performs tokenization, part-of-speech tagging, lemmatization, and dependency parsing according to the Universal Dependencies framework (UD),<sup>5</sup> and NameTag 3, a tool for named entity recognition.<sup>6</sup> In the Universal Dependencies framework, each token in a sentence is assigned a part-of-speech tag and represented as a node in a dependency tree.

Instead of attempting to simplify texts automatically (with the danger of altering the semantics), PONK functions as a diagnostic aid that signals segments that potentially deserve closer attention during the process of their drafting. The final decision on whether to modify the indicated passages is up to the author.

## 4 Implementation

The module runs as a server, and may be called via REST API.<sup>7</sup> We implemented the rules using the Udapi Python library.<sup>8</sup>

The input is supplied in the CoNNL-U format.<sup>9</sup> The output consists of the same CoNNL-U file with rule annotations in the MISC column, and of meta-information about the applied rules (such as rule names or user hints).

We stick to a few technical principles across all the rules. Many rules require counting (e.g. of the number of tokens in a given span). However, we found that simply counting the number of tokens may become misleading as the cohesion of constituents that tokens can combine into varies, especially in relation to grammatical agreement

between nouns and their attribute(s) or determinant(s).

To illustrate this decision, we propose two sentences (A) and (B). In both sentences, the object and its governing verb (both marked by bold font) are separated by eight tokens but the respective syntactic structures are of different complexity levels.

(A) “*Stanovený **rozsah**<sub>OBJECT-ACC</sub> omezení<sub>GEN</sub> má opatrovník<sub>NOM</sub> povinnost<sub>ACC</sub> při výkonu<sub>LOC</sub> opatrovnictví<sub>GEN</sub> **reflektovat** [...]*” (approx.: *The guardian has a duty to reflect the established scope of restrictions during the exercise of guardianship.*)  
 (B) “*není zřejmé, [...] **koho**<sub>OBJECT-ACC</sub> konkrétně předmětná<sub>NOM-ADJ</sub> stavba<sub>NOM</sub> svými<sub>INS-DET</sub> závadným<sub>INS-ADJ</sub> technickým<sub>INS-ADJ</sub> stavem<sub>INS</sub> ohrožuje.*” (approx.: *It is not clear [...] exactly whom the subject building endangers with its defective technical condition.*) In (A), apart from the verb “*má*” (*has*), which is part of the verbo-nominal predicate “*mít povinnost*”, and the preposition “*při*” (*during*), the two tokens are separated by five nouns in four different grammatical cases (cf. the subscripts), indicating a very diverse syntactic structure. On the other hand, in (B) the two tokens are separated by one adverb and two nominal phrases, each containing a single noun preceded by one to three attributes that match the noun’s grammatical case,<sup>10</sup> forming a less diverse, more easy-to-navigate syntactic structure.

We therefore consider adjectives and determinants to form single units with their governing tokens, which they are in morphological agreement with. Furthermore, we ignore punctuation, prepositions, and conjunctions, as they only mark syntactic relations. We call the resulting units “phrases”, as they serve as a proxy for phrases with comparable perceived difficulties during parsing. The phrases are used to determine any distance, position, length, or ratio referred to in our rules (see Section 5), unless otherwise specified.

In addition, we decided to exclude citations (marked by quotation marks) from the rules’ scopes, as the user will likely want to keep them intact. The module recognizes quotations by matching two quotation marks.

In this study, we call segments of texts indicated by the rules “violations”, as they violate the good practices in legal texts drafting.

<sup>4</sup>For more details, see: <https://github.com/ufal/ponk>

<sup>5</sup><https://lindat.mff.cuni.cz/services/udpipe/>, Straka (2018).

<sup>6</sup><https://lindat.mff.cuni.cz/services/nametag/>, (Straková and Straka, 2025).

<sup>7</sup>The source code is available at <https://github.com/ufal/ponk-linguistic-rules>.

<sup>8</sup><https://udapi.github.io/>

<sup>9</sup><https://universaldependencies.org/docs/format.html>

<sup>10</sup>It should be mentioned for the sake of completeness that apart from grammatical case, this type of morphological agreement in Czech also involves grammatical number and gender.

## 5 Rules

In total, there are twenty rules implemented in PONK. In addition to the definition of the rules offered in this paper, the rules are also briefly described in the PONK’s user manual (see [Hladká et al. \(2025\)](#)).

### 5.1 Sources of rules

One of the sources of inspiration for the rules were scientific works addressing readability. [Šamánková and Kubíková \(2022\)](#) and [Šváb \(2023\)](#) target the legal and administrative domains and the desirable drafting style within them. [Sgall and Panevová \(2014\)](#) refer to the “good standards” of standard Czech, although not necessarily in the context of the legal domain.

### 5.2 Categories of rules

We divide the rules into four categories as presented in the following list. From the perspective of traditional linguistic layers, the rules are related to syntactic or lexical phenomena.

1. Deprecated constructions and expressions (see Section 6)
2. Syntactic malconstructions rules
  - (a) Position rules (see Section 7.1)
  - (b) Ambiguity rules (see Section 7.2)
  - (c) Cluster rules (see Section 7.3)

## 6 Discouraged constructions and expressions

Discouraged expressions refer to expressions that are vague or ambiguous or otherwise decrease the readability, brevity, or easy-to-navigate structure of the text. In total, eight rules belong to this category. Because most rules from this category are searching for selected phrases, exceptions for common legal terms are included in each rule if relevant.

1. ANAPHORICREFERENCES: This rule concerns vague references to anaphoric elements, which may cause uncertainty and need to examine the preceding text to identify the referent. Examples include “Z výše uvedeného je zřejmé . . .” (*From the above mentioned, it can be concluded . . .*) and “S ohledem na tuto skutečnost . . .” (*With regard to this fact . . .*).

2. ABSTRACTNOUNS: Semantically vague nouns are often used as placeholders for a more precise expression. Common examples include

“podstata” (*nature*), “aspekt” (*aspect*), “stupeň” (*instance*), “okolnosti” (*circumstances*), etc. The nouns are contextually dependent, and they may occur in texts in a non-vague sense or as parts of legal terms. For example, *stupeň* may occur as part of the term “soud prvního stupně” (*court of first instance*).

3. CONFIRMATIONEXPRESSIONS: We define confirmation expressions as expressions that, instead of supposedly increasing certainty, paradoxically contribute to ambiguity. Examples include “nepochybně” (*without doubts*) and “rozhodně” (*surely*). On the other hand, expressions like “jednoznačně” (*unambiguously*) are not included because they do, in fact, increase the certainty of interpretation.

4. PASSIVEVOICE: Passive voice is grammatically correct and acceptable in Czech. However, it is considered slightly archaic and more difficult to navigate for readers, compared to active voice [Šamánková and Kubíková \(2022\)](#), which is recommended. That being said, passive voice is useful for sentences with unknown agent or for manipulating word order for the purposes of functional sentence perspective [Firbas \(2009\)](#). We therefore choose to only include passive constructions with an overt agent.

5. REDUNDANTEXPRESSIONS: Redundant expressions are expressions that have little to no function in text and can be omitted without loss or alteration of semantic content. Common examples include “v neposlední řadě” (*last but not least*), “je nutné zdůraznit” (*it is necessary to emphasize*), or “v kontextu věci” (*in the context of the matter*).

6. RELATIVISTICEXPRESSIONS: This rule concerns expressions that increase uncertainty. Examples include lexical units such as “snad” (*perhaps*), “jaksí” (*somehow*), “obdobně” (*similarly*), “jevit” (*to seem*), etc. We are aware that the ambiguity of some lexical units (e.g., “velmi”, (*very*) is contextually dependent. For this reason, we focus on expressions that are inherently uncertain.

7. TOOLONGEXPRESSIONS: This rule identifies expressions that can be replaced with synonymous, but shorter equivalents to ensure the conciseness of the text. The rule searches for several specific expressions, including “v důsledku toho”, “v případě, že”, or “za situace”, which can be replaced by “proto” (*therefore*), “pokud” (*if*), and “když” (*when*), respectively.

8. WEAKMEANINGWORDS: This rule signals lexical units (mostly verbs) that are considered se-

mentally vague and are often used as fillers in legal texts. Examples include “*zdát se*” (*to seem*) or “*ovlivnit*” (*to influence*). Users are encouraged to use more precise verbs instead. Let us note that the semantically vague verbs do not include verbs that participate in verbo-nominal predicates.<sup>11</sup>

## 7 Syntactic malconstructions rules

Suboptimal syntactic constructions, although grammatically correct, decrease the readability of the texts. Based on the nature of the constructions and the caused effect, we divide the rules into three subcategories: position rules, ambiguity rules, and cluster rules. The numeric thresholds relevant to individual rules were set after repeated experiments.

### 7.1 Position rules

The five position rules concern the distance between tokens with specific syntactic categories. Greater distance may make a sentence difficult to navigate, and a full understanding may require several re-reads (Šamánková and Kubíková, 2022). The position or distance of the relevant token(s) is established according to the method described in Section 4.

1. **PREDSUBJDISTANCE**: This rule addresses the distance between the predicate and the subject. The limit distance is 6. Violations of this rule may be triggered, among others, by adnominal clauses. In case of analytic predicate (verb forms including auxiliary verbs, modal verbs, or a copula; it is not to be confused with verbo-nominal predicate), the position of the auxiliary, copula, or the governing modal verb is taken into consideration, as they enter into agreement with the subject in Czech grammar.

2. **PREDOBJDISTANCE**: The limit distance between the predicate and the object is set to 6. In case of analytic predicates (see **PREDSUBJDISTANCE** above), we consider the position of the content token (e.g., “*byl proveden*”, *was conducted*), as the object belongs to its valency.

3. **MULTIPARTVERBS**: This rule addresses the distance between the parts of multipart verbs forms (passive voice, future tense, past tense, conditional mood). Reflexive verbs are not included, because the position of the clitic “*se*” is fairly restricted in Czech grammar and thus mostly not pertinent for our purpose. Constructions with modal verbs

are covered by the rule **INFVERBDISTANCE** (see below). The limit distance is 5.

4. **INFVERBDISTANCE**: The distance between verb and a dependent infinitive may concern two types of units: (1) constructions with modal verbs, and (2) infinitives that are objects dependent on a non-modal verb. The default limit distance is 5.

5. **PREDTOOFARINCLAUSE**: Although the dominant word order of Czech is SVO, it is possible (and grammatically correct) to create sentences with alternative word orders (including positioning the subject or the verbal predicate at the end of a sentence). According to the approach of functional sentence perspective (see Firbas (2009); Jasinskaja and Šimík (2023)), the non-dominant word orders are typically related to the information structure of the sentence. Due to declension, misinterpretations of such sentences are rare (although they may lead to temporary garden-path situations, Ceháková and Chromý (2023)).

As for the rule **PREDTOOFARINCLAUSE**, the tolerated limit position of a predicate is the 9th position from the beginning of a clause (not a whole sentence). We attempt to exclude<sup>12</sup> sentences with verbs at the end of a clause, since such word order can be interpreted as a likely conscious choice of putting the verb in focus (Jasinskaja and Šimík, 2023). The violations often occur in sentences with complex adjuncts.

### 7.2 Ambiguity rules

Ambiguity rules refer to situations in which suboptimal syntactic structures result in ambiguous constructions and decreased readability. Currently, there are two ambiguity rules.

1. **DOUBLEADPOS**: This rule addresses syntactic structures in which there are at least two coordinated tokens, of which the first is preceded by an overt preposition and the latter is not, leading to a possible ambiguity as for syntactic relations. An example may be “*OSPOD však upozorňují, že by se měl zabývat oznámeními otce o možné manipulaci dětí ze strany prarodičů, případně i jiných zásazích prarodičů do jeho výchovy.*” The absence of repetition of “*o*” leads to two potential interpretations: Either *[... ]reports about manipulation of the children by their grandparents or about other interferences by the grandparents.* or *[... ]reports about*

<sup>11</sup>For a more detailed analysis of verbo-nominal predicates in Czech, see Radimský (2017).

<sup>12</sup>The exclusion is currently formalized as the simultaneous occurrence of a predicate appearing with the last 3 tokens of a clause and being preceded by all its core arguments (including an infinitive).

*manipulation of the children by their grandparents or by other interferences by the grandparents.* If the distance between the two tokens is 4 or less, the rule is not triggered.

2. **INCOMPLETECONSTRUCTION**: This rule captures missing parts of multi-word constructions. If a part is omitted without the reader’s knowledge, it may lead to confusion or force the reader to re-read sections of the text, as they still expect the missing part to appear later. Examples include “*jednak ... jednak*” (both ... and), “*bud ... nebo*” (either ... or), and “*zaprvé ... zadruhé*” (firstly ... secondly).

### 7.3 Cluster rules

The five rules in this subcategory concern clusters of phenomena within short spans (usually a sentence or part of a sentence) of text. For further details on how the lengths of sentences or ratios of phenomena are counted, see Section 4.

1. **CASEREpetition** It is possible to encounter sequences in which many tokens share the same morphological case, but not as a result of morphological agreement. Consequently, parsing may require more cognitive load on the readers. The typical case that leads to violations of this rule is genitive, which is used for noun adjuncts in Czech.<sup>13</sup> The rule focuses on nouns and noun adjuncts ignores other parts of speech like adjectives or pronouns, which we consider to be easy to navigate in the text. The maximum tolerated number of nouns sharing the same case is 4. Appositions and coordinations are excluded.

2. **LONGSENTENCES** This rule examines the length of a sentence. The maximum length allowed is 22. The recommended modification of long sentences is to divide them into several shorter ones.

3. **TOOFEWVERBS** In this rule, the ratio of verbs to the length of a clause is analysed. The proportion considered as insufficient is 10% or higher. For this rule, every verbal lexeme (including modal verbs) is counted as one unit, regardless of whether the verb is complex or not.

4. **TOOMANYNEGATIONS** In Czech, negation concerns especially verbs, nouns, adjectives, and adverbs can be subjected to negation by the prefix *ne-*.<sup>14</sup> Multiple successive negations are generally

<sup>13</sup>In the traditional terminology of Czech grammar, we call the syntactic function “*přívlastek neshodný*.”

<sup>14</sup>Tokens belonging to other parts of speech are generally not morphologically negated, with several exceptions (some pronouns, particles, etc.).

discouraged (albeit grammatically correct) as they decrease readability. The rule is applied if at least 3 negations are present in a text span, and the left-most and right-most sentences of the span must contain at least 2 negations at the same time. Legal terms containing *ne-* are excluded from the rule: “*nezletil*” (minors), “*nevinný*” (not guilty), “*nedbalost*” (negligence), etc.

5. **TOOMANYNOMINALCONSTRUCTIONS** The proportion of nouns in a clause should not be greater than 45%. Coordinated nouns are excluded, so that enumerations do not trigger the rule. The rule also ignores named entities (conciseness is unlikely) and abbreviations typical for Czech legal texts (e.g., “*odst.*”, *paragraph*).

## 8 Rules application on authentic texts

We tested how successfully rules violations are identified. Furthermore, we examined whether a hypothetical modification of the identified segments leads to an increase in readability. We analyzed a corpus of authentic legal texts, compiled for the purpose of this research. Then, a human annotator evaluated each identified violation.

### 8.1 Corpus

We established a corpus of 30 texts (a subset of the KUK corpus)<sup>15</sup> produced by ombudsmen, the public defenders of rights in the Czech Republic. The texts cover a period from 2007 to 2023, with the majority being published in the 2010’s (18 texts). Each of the texts was assigned a document ID. The length of the individual texts varies.<sup>16</sup> The texts are written by different ombudsmen (although an ombudsman may author multiple documents constituting the corpus) and cover various themes, such as children’s homes, disputes in medical or education domains, social security benefits, and more. The texts were not edited or modified for the purpose of our research.

### 8.2 Annotation

The manual annotation was performed by a linguist with experience in legal texts and terminology. The annotator controlled all identified violations and classified them as *useful*, *not useful*, or *incorrect*.

<sup>15</sup><https://lindat.mff.cuni.cz/repository/items/929f7d3d-4783-4d8b-8378-ccd2f6014493>

<sup>16</sup>The longest text has 18,090 tokens (almost 110,000 characters without spaces), while the shortest text contains only 259 tokens (1,654 characters without spaces.)

The category *incorrect* indicates falsely positive violations that do not correspond to definition of the given rule. For example, in the case of the rule `INCOMPLETECONSTRUCTION`, there are identified spans of texts with the construction “*bud’ . . . nebo*”, which indicates that the second part “*nebo*” was omitted from the text, while in fact the second part is present but was not recognized by the system.<sup>17</sup> For this reason, such violations are evaluated as *incorrect*. Another example of incorrectly signaled violations concerns the rule `LONGSENTENCES`, as there are cases in which the system treated two successive sentences as a single sentence.

Both *useful* and *not useful* are used for correctly identified violations (true positives). The choice between them depends on whether the alteration of the indicated segment would contribute to a higher level of readability or not. The category *not useful* thus allows for a careful consideration of the context of each violation and examination of its impact. Correct but *not useful* violations typically consist of terminological units.<sup>18</sup> For example, the token “*podstata*” (*nature*) is included in the rule `ABSTRACTNOUNS`. However, it is part of the term “*konkursní podstata*” (approx.: *bankruptcy estate*) and for this reason it is evaluated as *not useful*, as any modification would alter the legal term. *Not useful* is also applied to any identified violation within a direct citation (see Section 4).

### 8.3 Results

With regard to the manual annotation, we can formulate several conclusions.

First, the frequency of violations of individual rules differs according to the character of the rule. As observed in Figure 2, `LONGSENTENCES` tracks the highest number of identified violations, followed by `TOOFEWVERBS`, `PREDSUBJDISTANCE`, and `ABSTRACTNOUNS`. The frequency of these violations seems to reflect the general perception of legal texts as texts of complex and hard-to-navigate syntactic structure and a certain degree of vagueness. By contrast, violations of `INCOMPLETECON-`

<sup>17</sup>An example from our corpus is: “Doporučuje se *bud’* ve VŘ podrobněji popsat pravidla pro vyřizování stížností (kdo, v jaké lhůtě, jakým způsobem), *nebo* takovou informaci vydanou ředitelem ústavu vyvěsit například na nástěnce oddělení.” Verbatim in English: “It is recommended to *either* describe in detail the rules for processing complaints [. . . ], or post the information issued by the office’s director on the office’s notice board.”

<sup>18</sup>Terminological units are generally excluded from the analysis, but as the exceptions are added manually, any not-yet-included term may be indicated by the system.

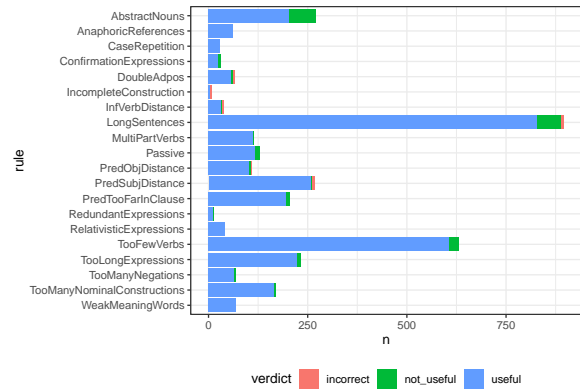


Figure 2: Overall frequencies of violations.

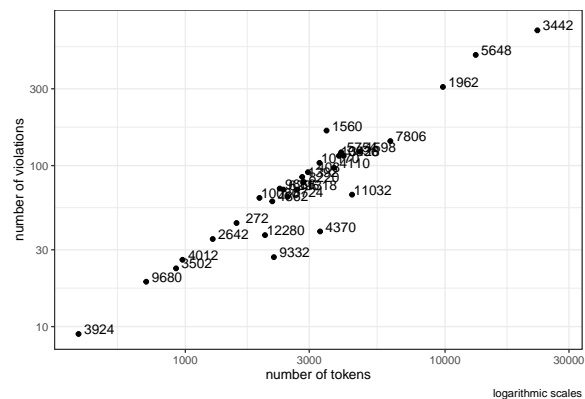


Figure 3: Number of all violations per document over document length. Document IDs are shown as labels.

`CONSTRUCTION` show the lowest absolute frequency, which is somewhat expected given the character and functions of the covered units. Perhaps curiously, violations `REDUNDANTEXPRESSIONS` and `RELATIVISTICEXPRESSIONS` also have relatively low frequencies, although we might expect them to occur more frequently in legal texts, as they both contribute to ambiguity and complexity. The low frequency may be influenced by the individual texts in the corpus or by the lexical units included in these two rules. More testing would be required to verify either hypothesis.

We can also state that the length of a document correlates with the frequency of violations (see Figure 3). As can be seen, text no. 3442 (18,090 tokens) contains a high number of violations across different rules, followed by texts no. 5648 (10,664 tokens) and no. 1962 (7,605 tokens). Observed minor deviations from this trend are connected to particularities of individual texts or influenced by the personal style of the author of the given text or by the topic of the text (especially with regard to not-yet-excluded terminological units).

In terms of the ratios of *useful*, *not useful*, and *incorrect* violations, the results are presented in Figure 4. As can be seen, *useful* indications dominate and exceed 75% in the case of all rules except INCOMPLETECONSTRUCTION. The high proportional rate of *incorrect* violations regarding this rule is due to several false positives, in which the system failed to recognize the second part of the multipart construction. Given the low absolute number of violations (7, out of which 3 are *incorrect*), we cannot identify the cause of false positives with certainty. We suspect an inaccurate text segmentation, or it may be anecdotal, but we have yet to verify this hypothesis. In the case of DOUBLEADPOS and INFVERBDISTANCE, the *incorrect* violations are caused mainly by false signaling of different parts of speech (such as verbal nouns in the case of INFVERBDISTANCE). We suspect that these false positives may be caused by the tokenization process performed by UDPipe but this remains to be verified.

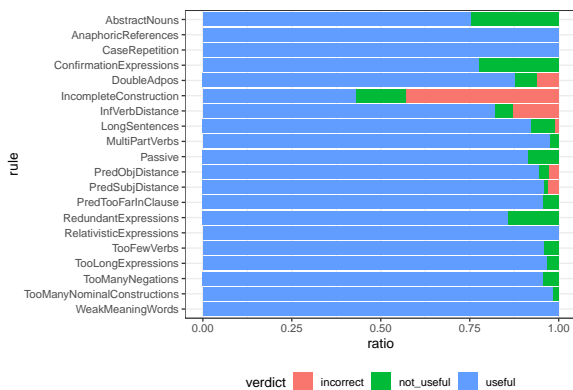


Figure 4: Ratios of useful, not useful, and incorrect violations.

The relatively high ratio of *non-useful* violations of ABSTRACTNOUNS is caused by terminological units that contain a trigger expression. These terms are to be included in the system as exceptions. In addition to terms, the token “*jasně*” (*clearly*) will become an exception to the rule CONFIRMATION-EXPRESSIONS. Similarly, we identified “*v této situaci*” (*in this situation*) as a possible exception to the rule REDUNDANTEXPRESSIONS. As for the other *not useful* instances, virtually any rule’s violations can occur in direct citations (see Section 8.2 where we discussed this in detail).

## 9 Discussion

The results of our experiments are promising. Based on human evaluation, if the text segments indicated by our rules were modified, it might significantly increase the readability and clarity of the text and reduce ambiguity and vagueness. Thus, the PONK tool seems to provide useful feedback to users and guides them towards good practices in legal or administrative text drafting in Czech.

However, some rules still exhibit a non-trivial amount of *not useful* identified violations. This is mostly the case for the discouraged-constructions rules, as they focus on hand-picked phrases or their variations. But since meaning is generally contextually dependent and our approach does not allow us to adequately represent semantics (we only operate with lemmatization and syntactic parsing), it remains challenging to identify only phrases that truly violate the good practices while excluding “non-problematic” ones. Similarly, a considerable number of *not useful* violations is due to their terminological character or participation in direct citations. Although we tried to exclude such situations from our analysis, the results show that there is still room for improvement and further refinement.

Despite the fact that not all identified violations were evaluated as useful, we do not assume that it would significantly hinder user experience. In addition, recommendations for each violated rule and user’s manual with more detailed definitions and examples are available in the PONK UI, providing explanations and a guide for potentially contentious or controversial violations.

## 10 Conclusions

In this paper, we presented a rule-based system, supported by a web application, that serves as an advisory tool called PONK, designed for authors of legal texts to help them increase the clarity and readability of texts. The system is adapted to Czech language and its grammar.

Following an experiment in which the system identifies phenomena non-conform to the established rules set (called “violations”) in authentic Czech legal texts and human annotation of the violations, we can formulate several conclusions based on the results. The majority of the rules show high percentage (over 75%, some of the rules up to 100%) of correctly indicated violations that, if eliminated, would contribute to a more readable legal text. That being said, the experiment revealed

that some of the rules are not applied by the system as expected, and thus more testing on a larger corpus and modifications are necessary. Similarly, a thorough analysis of legal terminology that could potentially trigger the rules is required.

Despite these obstacles, we consider the system’s overall performance satisfactory and positively contributing to the readability of Czech legal text in two main aspects. First, it will direct users’ attention towards the parts of texts that contain characteristics of legalese or suboptimal linguistic expressions that may distract the reader. And second, providing recommendations on how to improve such segments of text might positively contribute to the readability of the users’ future texts, possibly outreaching the selection of phenomena the rules are designed to capture.

Regarding future research in this area, several more rules are currently being considered for implementation. We also experimented with deterministic repair suggestions for multiple rules. However, for most rules, the applied tools lacked sufficient understanding of the linguistic structure (especially semantics and pragmatics), which would be crucial for offering reliable suggestions. The current state of the PONK tool does not include Large Language Models (LLMs) due to (1) their non-determinism and (2) privacy concerns. However, the potential positive impact of a hybrid system that includes LLMs guided by readability principles is something we are currently exploring, combining the linguistic precision of the PONK rules with the flexibility and creativity of LLMs.

## Limitations

Although our rules are applied exclusively to legal texts, the vast majority of the targeted linguistic phenomena are not unique to the legal domain. Syntactic complexity, vague expressions, or structurally dense constructions can occur in virtually any Czech text, written or spoken. However, they tend to be particularly prominent in legal and administrative writing, where traditional drafting conventions often favor formal and complex formulations. The resulting misunderstandings, ambiguity, or reduced readability can directly affect the interpretation of legal texts, which underscores the importance of addressing linguistic complexity in this domain.

Our presented paper is focused on the Czech legal language and the examined phenomena are

closely tied to specific grammatical and stylistic properties of Czech. For this reason, the rules cannot be transferred directly to other languages without modification. Nevertheless, we believe that the proposed rule set and the underlying methodological approach may serve as a useful starting point for similar systems developed for other languages. Some rules would likely require only minor adjustments, while others might be less applicable in languages with different grammatical structures (e.g., rules related to declension such as CASEREPE- TITION).

Since we focused on designing rules that provide useful feedback, we may have made some of them too strict to detect all violations of good readability practices targeted by the rules. Our evaluation design is not able to reflect on this, as it only explores the correctness or usefulness of the rules after they have been applied but not if they are applied in all relevant cases. However, we believe that the users’ interaction with our module may still make them proactively notice such violations when writing other texts in the future.

Another limitation of the presented research is that it does not address legal terminology. Specialized legal terms undoubtedly contribute to the perceived opacity of legal texts, particularly for readers without legal training. Such readers may not know the precise definitions of such terms or may interpret them only through their generalized meanings in everyday language, as noted by Cornu (2005). At the same time, legal terminology constitutes an inherent part of legal discourse and therefore cannot simply be rewritten or replaced without affecting legal precision. For this reason, our present work focuses on other linguistic aspects of readability. Nevertheless, the question of accessibility of legal terms to non-expert readers remains an important direction for future research.

Our objective is not to create another normative grammar handbook, but rather to contribute to the development of tools that promote clearer and more reader-friendly legal and administrative writing. From this perspective, any step that improves the accessibility of legal language for the general public is both meaningful and necessary.

## Acknowledgements

This research received financial support by the projects: “PONK - Client-Oriented Writing” (TAČR Sigma, No. TQ01000526), “Human-

centred AI for a Sustainable and Adaptive Society” (No. CZ.02.01.01/00/23\_025/0008691, co-funded by the European Union and LINDAT/CLARIAH-CZ (No. LM2023062, supported by the Ministry of Education, Youth and Sports of the Czech Republic).

The work described herein has been using services provided by the LINDAT/CLARIAH-CZ Research Infrastructure (<https://lindat.cz>).

## References

- Kurt Baldinger. 1984. *Vers une sémantique moderne*. Klincksieck.
- Libuše Halfarová Bohuslav Halfar and Michaela Bučková. 2022. *Linguistic and legal comprehensibility of the text*. *Ekonomická revue - Central European Review of Economic*.
- Markéta Ceháková and Jan Chromý. 2023. Garden-path sentences and the diversity of their (mis)representations. *PLoS ONE*, 18(7).
- Robert P. Charrow and Veda R. Charrow. 1979. Making legal language understandable: A psycholinguistic study of jury instructions. *Columbia Law Review*.
- Marta Chromá. 2016. *Právní překlad v teorii a praxi. Nový občanský zákoník*. Karolinum.
- Jan Chromý, Silvie Cinková, and Jana Šamánková. 2021. Srozumitelnost českého odborného a úředního textu — proč se jí zabývat a jak ji měřit. *Studie z aplikované lingvistiky*, 12(1):38–52.
- Gérard Cornu. 2005. *Linguistique juridique, 3e édition*. Paris : Montchrestien.
- Shari Seidman Diamond, Beth Murphy, and Mary R Rose. 2012. The kettleful of law in real jury deliberations: Successes, failures, and next steps. *Nw. UL Rev.*, 106:1537.
- Jan Firbas. 2009. *Functional Sentence Perspective in Written and Spoken Communication*. Cambridge University Press.
- Yu Han, Aaron Ceross, and Jeroen H. M. Bergmann. 2024. *The Use of Readability Metrics in Legal Text: A Systematic Literature Review*. *arXiv preprint*. ArXiv:2411.09497 [cs].
- Alissa J. Hartig and Xiaofei Lu. 2014. *Plain English and legal writing: Comparing expert and novice writers*. *English for Specific Purposes*, 33:87–96.
- Barbora Hladká, Silvie Cinková, Jan Černý, Vítek Eichler, Tomáš Knap, Ivan Kraus, Barbora Kubíková, Ivana Kvapilíková, Jiří Mírovský, Tereza Novotná, Tomáš Polák, Arnold Stanovský, Jana Šamánková, Michal Kuk, and Přemysl Pospíšil. 2025. *Srozumitelnost českých právních a administrativních dokumentů ve výzkumu a praxi*. Technical Report TR-2025-75, Prague, Czech Republic.
- Kateřina Hodková. 2020. Romanisms in the Czech New Civil Code and their French Equivalents. *Acta Faculty filozofické Západočeské univerzity v Plzni*, 12(1):61–82.
- Kateřina Hodková. 2022. Les relation sémantiques au carrefour des champs conceptuels du droit tchèque et du droit français. *Studia Romanistica*, 22(1).
- Kateřina Hodková, Jana Pešková, Ivo Petrů, and Jan Radimský. 2021. Kontrastivní srovnávání právní terminologie na příkladech z vícejazyčné databáze právních termínů legterm. *Jazyk a kultura*.
- Katja Jasinskaja and Radek Šimík. 2023. *Slavonic free word order*. To appear in *The Oxford guide to Slavonic languages* (eds. Jan Fellerer & Neil Bermel).
- Ivan Kraus. 2025. *Predicting readability of czech legal writing using linguistic features*.
- Eric Martínez, Francis Mollica, and Edward Gibson. 2022. *Poor writing, not specialized concepts, drives processing difficulty in legal language*. *Cognition*, 224.
- Tereza Novotná, Jan Černý, Ivan Kraus, Ivana Kvapilíková, Jiří Mírovský, Arnold Stanovský, and Barbora Hladká. 2025. *PONK: Tool for Client-Oriented Legal Writing in Czech*. In *JURIX 2025: The Thirty-eighth Annual Conference*, volume 416 of *Frontiers in Artificial Intelligence and Applications*, pages 330–336, Amsterdam, Netherlands. University of Turin, Italy, IOS Press.
- Jan Radimský. 2017. *Analytický predikát s kategoriálním slovesem*. *Nový encyklopedický slovník češtiny*.
- Petr Sgall and Jarmila Panevová. 2014. *Jak psát a jak nepsat česky*. Karolinum.
- Milan Straka. 2018. *UDPipe 2.0 Prototype at CoNLL 2018 UD Shared Task*. *Proceedings of the CoNLL 2018 Shared Task: Multilingual Parsing from Raw Text to Universal Dependencies*.
- Jana Straková and Milan Straka. 2025. *NameTag 3: A tool and a service for multilingual/multitagset NER*. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 3: System Demonstrations)*, pages 31–39, Vienna, Austria. Association for Computational Linguistics.
- The GDS Team. 2014. *Guest post: Clarity is king – the evidence that reveals the desperate need to re-think the way we write*. *Government Digital Service*.
- Michal Tomášek. 2003. *Překlad v právní praxi, 2nd edition*. Linde.

Veřejný ochránce práv. 2024. [Testování srozumitelnosti](#). Výzkumná zpráva KVOP-17384/2024, Veřejný ochránce práv, Brno.

Jana Šamánková and Barbora Kubíková, editors. 2022. *Jak psát srozumitelné úřední texty*. Kancelář veřejného ochránce práv.

Jakub Šváb. 2023. *Jak psát, aby se to dalo číst*. Leges.

# Human–AI Annotation Error Auditing for Hebrew Diacritization with Frontier LLMs

Hillel Gershuni and Avi Shmidman

Bar-Ilan University and DICTA  
gershuni@gmail.com, avi.shmidman@biu.ac.il

## Abstract

Large annotated datasets inevitably contain errors that are costly to identify via manual review. We study a human–AI annotation error auditing workflow using frontier Large Language Models (LLMs), focusing on Hebrew *nikud* (diacritization). We take the the EACL 2023 Hebrew Homograph Challenge Set as our test case. In a focused evaluation on 12 of the homograph sets with 271 confirmed errors (verified through exhaustive manual review of all 7,241 sentences), Gemini 3 Pro achieves 83.6% recall (95% confidence interval: [79.3%, 88.2%]) and 99.1% precision - substantially higher than other frontier LLMs. Two independent human experts achieved 62.4% and 42.8% recall respectively, a 20-percentage-point spread that reflects the difficulty of sparse-target error search. Even the union of both experts’ findings (73.4% recall) falls short of a single LLM run (83.6%), while LLM-aided auditing reduces review effort by over 95%. We analyze the trade-offs between batch size and recall, and release both a human-verified Gold Standard with per-error difficulty annotations and a globally corrected version of the Challenge Set.

## 1 Introduction

High-quality annotations are foundational for evaluating and training NLP systems, but manual error auditing is slow and expensive. This challenge is particularly acute in the case of annotated datasets of Hebrew *nikud* (diacritization). Hebrew is typically written without vowel diacritics; in an annotated Hebrew corpus for *nikud*, the words are annotated with the relevant diacritic vowels. However, the determination of the correct *nikud* is not a simple dictionary lookup; rather, the task often requires disambiguating homographs using deep contextual cues. Even expert-curated resources can accumulate subtle errors that impact downstream model reliability. A case in point is the EACL 2023 Hebrew Homograph Challenge Set (Shmidman et al.,

2023), upon which we will focus herein. This Challenge Set is comprised of ~150,000 sentences, each containing a Hebrew homograph which is annotated with diacritics to disambiguate its meaning. Naturally, in a human-curated dataset of this size, a certain number of errors will persist. The question is how to locate them.

We propose using frontier LLMs as a human–AI auditing tool to flag likely *nikud* errors. Unlike the traditional diacritization scenario, in which computational systems are built to diacritize non-diacritized texts (see below, 3.1), our auditing approach focuses on validating *existing* diacritization labels with a prioritized **recall-centric objective**, surfacing as many real issues as possible for subsequent human verification.

Our key contributions are as follows:

- We define a practical annotation error auditing procedure and quantify precision/recall trade-offs across multiple frontier LLMs.
- We quantify the impact of batch size on error detection sensitivity, identifying a practical sweet spot that balances recall and cost.
- We validate the pipeline through a workflow comparison with two independent human experts, reporting inter-annotator agreement ( $\kappa = 0.57$ ) and showing that even the union of both experts’ findings (73.4% recall) falls short of a single LLM run (83.6%), while LLM-aided auditing reduces review effort by over 95%.
- We provide a human-verified gold standard covering 7,241 sentences across 15 homograph files (12 homograph groups), with 271 confirmed diacritization errors and per-error difficulty annotations derived from cross-system detection rates, which can be used as a benchmark for future experiments in this area.

- We apply our pipeline to the entirety of the aforementioned EACL 2023 Hebrew Homograph Challenge Set, and we release a new globally corrected version of this Challenge Set for the Hebrew NLP community.<sup>1</sup>

## 2 Task and Linguistic Background

The primary challenge in Hebrew diacritization auditing lies in the morphological density of the language. A single consonantal string can represent multiple parts of speech, tenses, or complex syntactic constructions involving suffixes.

**Motivating Example.** Consider the string  $\text{אהבה}$  (/ʔhbbh/). Depending on the context, it may be diacritized as:

1. **Noun:**  $\text{אהבה}$  (love, /ʔa.hă.va:/).
2. **Verb (3fs Past):**  $\text{אהבה}$  (she loved, /ʔa:.hă.va:/).
3. **Verb + Accusative Suffix:**  $\text{אהבה}$  (he loved her, /ʔă.he.'va:h/).<sup>2</sup>

In these examples, the *nikud* distinction, for example between *patah* (ֶ, historically [a]), *qamatz* (ֵ, historically [a:]), and *hataf-patah* (ֶּ, historically [ă]), is contextually governed by gender agreement or the presence of an accusative pronominal suffix (marked by a *mapiq* dot in the final ם). While the Tiberian diacritization system originally distinguished the quality and length of these vowels (with *qamatz* representing an intermediate [ɔ] sound), Modern Hebrew follows a Sephardic reading tradition in which these distinctions have merged as [a] in the relevant contexts. The *mapiq*, too, is consistently unpronounced in standard speech. Human auditors therefore cannot rely on phonetic intuition, as the orthography preserves historical-grammatical layers that have no auditory counterpart in the modern language. This complexity makes manual review at scale prohibitively costly.

<sup>1</sup>Released artifacts are at [https://github.com/Dicta-Israel-Center-for-Text-Analysis/EACL\\_2023](https://github.com/Dicta-Israel-Center-for-Text-Analysis/EACL_2023). The original EACL 2023 Challenge Set is preserved under git tag v1.0; the corrected Challenge Set, Gold Standard, and difficulty annotations are released under tag v2.0.

<sup>2</sup>This particular form combines a verb with a cliticized pronominal suffix, and would in principle be resolvable through morphological segmentation. However, Hebrew digital texts carry no segmentation markup: the entire orthographic string is encountered as a single token, indistinguishable from the homographic noun and verb forms above. Determining the correct diacritization therefore requires precisely the kind of contextual analysis described in this paper, and the same is true for other homographs in our evaluation set.

## 3 Related Work

### 3.1 Hebrew diacritization.

Hebrew diacritic restoration has been approached using rule-based systems (Choueka and Neeman, 1995) and hidden markov models (Gal, 2002), as well as neural-network-based systems such as *Dicta Nakdan* (Shmidman et al., 2020) and *Nakdimon* (Gershuni and Pinter, 2022), which take consonantal texts as input and predict the diacritics. In contrast, our work targets auditing: given a diacritization label, we ask an LLM to judge its contextual correctness.

### 3.2 Annotation error detection.

Automated detection of labeling errors has been studied using training dynamics (Swayamdipta et al., 2020) or probabilistic methods (Northcutt et al., 2021). LLMs provide a complementary tool: they can draw on broad knowledge and generate rationales without task-specific training. Our work contributes an empirical comparison of frontier LLM models in a realistic human-in-the-loop workflow.

## 4 Methodology

### 4.1 Source Dataset

We utilize the EACL 2023 Hebrew Homograph Challenge Set (Shmidman et al., 2023), containing 75 homographs across ~150,000 sentences. The corpus consists exclusively of **written Modern Hebrew**, drawn from a mixture of newspapers, Wikipedia, literature, and social media. Per-homograph sentence counts (typically 1,000 sentences for the primary analysis and 250–500 for secondary analyses) are tabulated in Appendix A of Shmidman et al. (2023). All auditing experiments reported in this paper use the original release of the Challenge Set; the corrected version is released as an output of this work.

### 4.2 Request Structure and Context Enrichment

Every processing request, regardless of batch size, included a fixed header consisting of:

- **System Prompt:** A linguist-persona prompt prioritizing high recall of *nikud* errors.
- (optionally) **Morphological Data:** A list of valid diacritized forms, lemmas, and grammatical features for the target homograph.

- (optionally) **Dictionary Definitions:** Semantic definitions and common usage examples for each morphological variant.<sup>3</sup>

This prompt was prepended to every batch of 50, 100, 200, or 500 sentences, allowing us to measure how batch size affects model performance while amortizing the fixed prompt cost across more examples. The full API request is given in Appendix D).

### 4.3 Experimental Design

We curated a **Focused Evaluation Subset** of 15 files (7,241 sentences total, gathered from 12 homograph groups), from across the EACL 2023 Hebrew Homograph Challenge Set. The selection proceeded as follows: an initial exploratory phase screened 25 homograph files from the alphabetical extremes of the dataset (homographs beginning with א–ק and ק) using multiple LLMs to flag suspected errors. From the results, we selected homograph groups with the *fewest* flagged errors, i.e., those where exhaustive human adjudication was feasible, to ensure a rigorous, fully verified Gold Standard. This criterion naturally filtered out high-noise homographs such as אס (see §6.4), whose large volume of prescriptive-norm disputes would have made complete adjudication impractical. The selection criteria were fixed before any focused evaluation runs were conducted.

This subset represents a diverse range of linguistic challenges including part-of-speech ambiguity, construct-state vs. attributive distinctions and inflectional errors.

**Models and Parameters.** We evaluate five frontier LLMs: Gemini 3 Pro, Gemini 3 Flash, GPT-5.2, and Claude Opus 4.5/4.6, all accessed via OpenRouter APIs in a zero-shot setting with the same prompt template (Appendix D) and structured JSON output. All runs used default temperature ( $T=1.0$ ), following Google’s recommendations for reasoning-centric generation (Google, 2026); temperature sensitivity is left for future work. No system-level seed was set; stochastic variation between runs is part of our analysis. A total of 50 valid focused runs were conducted.<sup>4</sup> Full model

<sup>3</sup>Definitions are sourced from the Academy of the Hebrew Language: <https://hebrew-academy.org.il>

<sup>4</sup>This count includes 10 exploratory runs not analyzed individually: 3 with Grok 4.1 Fast (recall <22%), 1 with Claude Sonnet 4.5 (recall <4%), and 6 early Gemini Pro runs using alternative prompt variants. All contribute to the aggregate union and difficulty statistics reported in §7.2 and §4.4. Results from two Gemini-Pro runs were discarded due to an incorrect reasoning-mode configuration.

identifiers, access dates, and additional configuration details are provided in Appendix G.

### 4.4 Gold Standard (GS) Construction

The Gold Standard for the focused subset was established through a multi-stage process designed to maximize recall:

1. **Flag collection:** Candidate errors were flagged independently by multiple LLMs across different configurations and by two independent human experts (both native Hebrew speakers with graduate-level linguistics training), all without access to each other’s output.
2. **Union aggregation:** All unique flags were pooled into a candidate set.
3. **Expert adjudication:** The first author reviewed every candidate, classifying each as *Error*, *Correct*, or *Other*. This yielded 265 confirmed errors across 863 adjudicated sentences.
4. **Exhaustive dormant-error audit:** To guard against errors absent from the candidate set, the first author manually reviewed all 6,378 remaining sentences, uncovering 6 additional errors (dormant rate: 0.09%).
5. **Adjudication robustness check:** Each expert’s findings were compared against the final GS. Of 178 findings by Expert 1 and 121 by Expert 2, only 9 and 5 respectively targeted sentences that were labelled *Correct* in the Gold Standard (13 unique cases, as one sentence was flagged by both). All 13 were re-adjudicated; every Gold Standard label was confirmed, leaving the Gold Standard unchanged.

The final Gold Standard contains **271** confirmed errors across all 7,241 sentences; all recall figures below are computed against this fully verified denominator.

**Error Difficulty Annotations.** Each error is annotated with a difficulty label based on its cross-system detection rate (proportion of 50 runs that flagged it): *Easy* ( $\geq 80\%$ ; 117 errors), *Medium* (40–80%; 109), *Hard* ( $< 40\%$ ; 40), or *Undetected* (5). These labels are included in the released Gold Standard to support stratified evaluation.

## 5 Experiments and Results

Before reporting results on the focused subset, we conducted a sanity check using contaminated files with 308 synthetically injected errors (Appendix F); all three top models achieved  $\geq 97\%$  precision and  $\geq 79\%$  recall, confirming basic detection capability.<sup>5</sup>

Table 1 compares all models across three batch sizes. Consistent with our auditing objective, we prioritize **Recall** as the primary metric. Gemini 3 Pro achieves the highest recall at every batch size, with 83.0%–83.6% recall and  $\geq 99\%$  precision across batch sizes 50–200 (95% confidence interval for batch size 200: [79.3%, 88.2%]). All pairwise recall differences between Gemini 3 Pro and other models are statistically significant (McNemar’s test with continuity correction,  $p < 0.001$  for all comparisons at every batch size). GPT-5.2 ranks second with reasonable recall at batch size 50 (69.4%) but substantially lower precision (85.1%).

Model	Recall (%) $\uparrow$			Precision (%)		
	B50	B200	B500	B50	B200	B500
Gemini 3 Pro	83.0	83.6	72.1	99.6	99.1	99.5
GPT-5.2	69.4	55.2	43.7	85.1	94.9	98.3
Claude Opus 4	45.6	43.5	40.0	56.4	94.4	68.8
Gemini 3 Flash	59.0	45.2	27.1	39.7	40.4	57.9

Table 1: Model comparison on focused evaluation set (271 errors), no morphology, no dictionary (–m–d). Gemini 3 Pro values are averaged over 2 runs. B50, B200, etc. indicate batch sizes.

### 5.1 Batch Size Trends and Efficiency

As shown in Tables 1 and 3 and Figure 1, Gemini 3 Pro recall remains stable across batch sizes 50, 100, and 200 (83.0%–83.6% averaged across configurations); McNemar’s tests confirm no statistically significant differences among these three sizes ( $p > 0.8$ ). The 200-sentence batch thus represents a practical “sweet spot” for large-scale auditing: equivalent sensitivity at substantially lower cost (Table 2).

However, a **sharp recall drop** occurs at batch size 500, where average recall falls to 71.4% (–11 pp from batch size 200). This drop is highly significant (McNemar’s test,  $p < 0.001$  for all B200 vs. B500 run pairs). The cause is likely multifactorial: longer prompts, output length constraints, or

<sup>5</sup>Because injected errors are often more salient than naturally occurring mistakes, these results should be interpreted as a sanity check rather than a proxy for real-world recall.

reduced per-sentence attention. But the practical implication is clear: batches beyond 200 sentences incur a significant sensitivity penalty that outweighs cost savings.<sup>6</sup>

**Cost and Review Burden.** Table 2 summarizes the cost–recall trade-off. At batch size 200, a single Gemini 3 Pro run costs approximately \$7 for our 7,241-sentence evaluation set (\$0.96 per 1,000 sentences) and completes in  $\sim 75$  minutes of wall-clock time. In terms of *review burden*, i.e. the number of flagged findings a human reviewer must adjudicate, Gemini 3 Pro at batch size 200 produces  $\sim 225$  findings on average, yielding 31.0 findings per 1,000 sentences. This reduces the human task from scanning 7,241 sentences to verifying  $\sim 225$  items, a 96.9% reduction in review volume.

Batch	Rec.	Cost	\$/1K	Time	Flags/1K
B50	83.0%	\$12.67	\$1.75	139m	31.4
B100	83.6%	\$9.09	\$1.26	104m	31.6
B200	82.2%	\$6.96	\$0.96	75m	31.0
B500	71.4%	\$4.28	\$0.59	41m	26.7

Table 2: Cost–efficiency comparison for Gemini 3 Pro (averaged across configurations). \$/1K = cost per 1,000 sentences. Flags/1K = flagged findings per 1,000 sentences, representing the human review burden. B = batch size. B200 achieves comparable recall at 77% of B100 cost and 55% of B50 cost.

### 5.2 Ablation Study: Context Robustness and Semantic Noise

We conducted an ablation study on the auxiliary context features (Morphology and Dictionary) to understand their impact on model performance. As shown in Table 3, the results indicate that explicit morphological or dictionary data does not yield consistent improvements under our tested zero-shot conditions.

#### Auxiliary Context Does Not Help Consistently.

No configuration consistently outperforms the others across batch sizes (Table 3): the bare –m–d setting (neither morphology nor dictionary) wins

<sup>6</sup>Note that because source files rarely divide evenly by the batch size, the final batch of each file is typically shorter (e.g., 100–174 sentences at batch size 200). It would have been defensible to discard these tail batches altogether, since they do not strictly conform to the nominal batch size of the experimental condition. We chose to retain them, however, because (i) discarding them would have reduced the effective evaluation set and (ii) we verified empirically that they do not inflate recall: the mean detection rate for errors in these tail batches (67.9%) was within 2.2 pp of the rate in full-sized batches (65.7%). The effect is therefore negligible.

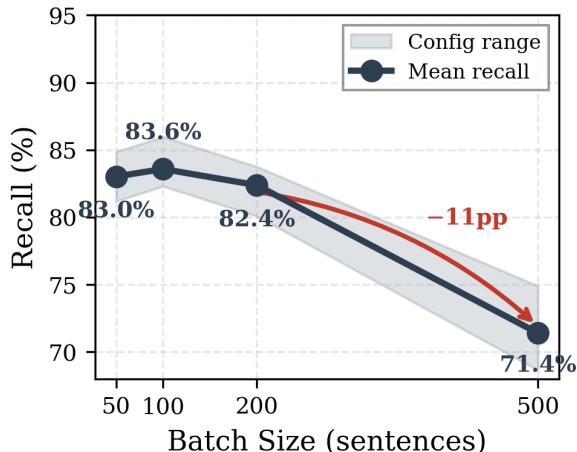


Figure 1: Recall vs. batch size for Gemini 3 Pro across all four context configurations (with and without morphological and dictionary info). Recall is stable from batch size 50 to 200, then drops sharply at 500. Context configuration has minimal effect relative to batch size.

at batch size 200, while +m-d (morphology only) wins at other batch sizes: 50, 100, and 500. The largest observed effect is at batch size 100, where +m-d achieves 233 TP vs. 223 for -m-d (+10 TP), though this is based on single runs and may reflect stochastic variation. McNemar’s tests confirm no statistically significant pairwise differences at batch sizes 50, 100, and 200 (all  $p > 0.28$ ); the inter-configuration gaps (5–10 TP) fall within run-to-run variance. These results suggest that the LLM possesses sufficient internal linguistic knowledge for disambiguation, and that it is not in need of extra morphological/dictionary info in order to make its determination.

Config	B50	B100	B200 <sup>†</sup>	B500
-m-d (bare)	225	223	<b>226.5</b>	194
-m+d (dict)	220	226	225.5	186
+m-d (morph)	<b>230</b>	<b>233</b>	221.5	<b>203</b>
+m+d (both)	225	223.5 <sup>†</sup>	217.5	184.5
<i>GT total</i>	271	271	271	271

Table 3: Ablation: True Positives across batch sizes for Gemini 3 Pro, varying whether morphological (m) or dictionary (d) data is provided within the prompt regarding the target homograph. B = batch size. (<sup>†</sup>Averaged over 2 runs per configuration.) Auxiliary context does not yield consistent improvement; at batch size 100, morphology shows the largest single-configuration effect (+10 TP), but this is based on single runs.

## 6 Error Analysis

The focused evaluation subset (15 files, 7,241 sentences) contains 271 confirmed errors, an error rate of 3.7%. Sections 6.1–6.2 analyze error patterns within this verified subset; §6.3–6.4 then discuss observations from the broader Challenge Set (332 files).

### 6.1 Syntactic Ambiguity: Construct vs. Adjective

Errors that were not found by most models are characterized by a high level of obscurity. Edge cases often involve the morphological distinction between construct-state nouns and attributive adjectives. In phrases like *התפרצויות אלימות*, the same consonantal string admits two competing analyses, distinguished only by the diacritization of the second word:

(a) *התפרצויות אלימות*  
*hitparšuyot 'alimut*  
 outburst.F.PL violence.F.SG  
 ‘outbursts of violence’ (N + N.CSTR)

(b) *התפרצויות אלימות*  
*hitparšuyot 'alimot*  
 outburst.F.PL violent.F.PL  
 ‘violent outbursts’ (N + ADJ)

Determining which reading is intended is a high-level syntactic challenge that can perplex even expert human annotators.

### 6.2 Errors Undetectable by LLMs

Of the 271 gold-standard errors, **five** (2%) were never flagged by any of the 50 valid runs. Two are construct-vs.-standalone ambiguities related to the syntactic challenges described in §6.1; the remaining three are dormant *אִשָּׁה* singular-vs.-plural cases uncovered only by exhaustive manual review (§4.4), that no model and neither human expert detected. Three additional dormant errors were each detected by only 1–2 runs, and seven more errors were each flagged by only a single run. All fifteen of these highly challenging cases share a common profile: the error is not clearly wrong in isolation, and disambiguation requires pragmatic or collocational knowledge that goes beyond syntactic parsing.

### 6.3 World Knowledge in Error Detection

Beyond the focused 15-file subset, exploratory runs on the broader Challenge Set (332 files) revealed error patterns that require encyclopedic, rather than linguistic, knowledge. Our analysis highlights a clear advantage of LLMs over rule-based or purely

Homograph	Context	Original	Correction	Error Type
אהבה /ʔhbh/	...גבר אחד «אהבה» רק גבר אחד... 'My mother <b>loved</b> only one man...'	אהבה (N) 'love' /ʔa.ha.'va/	אהבה (V) 'she loved' /ʔa.ha.'va/	Part-of-Speech (Noun → Verb)
אלימות /ʔlijmwt/	...ישירות ובלתי «אלימות»... '...direct and non- <b>violent</b> '	אלימות (N) 'violence' /ʔa.li.'mut/	אלימות (Adj) 'violent' /ʔa.li.'mot/	Syntax (Noun → Adj)
תנאי /tnʔj/	...בסיסי לחברה... 'a basic <b>condition</b> for society...'	תנאי (Pl.C) 'conditions of' /tna.'e/	תנאי (Sg) 'condition' /t.'nai/	Inflection (Pl → Sg)
בהי /bxjj/	...הקדושים «בהי»... 'in the <b>life of</b> the holy...'	בהי (Poss) 'in my life' /be.xa.'jai/	בהי (Cst) 'in the life of' /be.xa.'jej/	Morphology (Poss → Cst)
אמרה /ʔmrh/	...סטודמאיר הוא שחקן... ' <b>Amar'e</b> Stoudemire is a player...'	אמרה (N) 'saying' /ʔim.'ra/	אמרה Amar'e (name)	World Knowledge (Named Entity)

Table 4: Representative errors identified by the auditing pipeline. *Original* is the incorrect diacritization in the source dataset; *Correction* is the LLM’s suggestion. IPA transcriptions reflect Modern Hebrew pronunciation; note that the אהבה pair is **phonetically identical**; the distinction is purely orthographic (§2). N=Noun, V=Verb, Adj=Adjective, Pl=Plural, Pl.C=Plural Construct, Sg=Singular, Poss=Possessive, Cst=Construct.

morphological systems: the integration of **broad world knowledge**. This advantage is most apparent in the detection of transliterated names that function as homographs to common Hebrew words, a class of errors that is virtually invisible to traditional Hebrew diacritizers lacking encyclopedic knowledge.

**Named Entities as Homographs.** The homograph אמרה (/ʔmrh/; in Hebrew usually /am.'ra/ “she said” or /im.'ra/ “saying”) frequently appears in sentences about people whose names are transliterated with the same consonantal form: athletes, politicians, and artists such as Amar’e Stoudemire, Emre Belözoglu, and Fadela Amara (see Table 4, last row). In the homograph challenge set they were diacritized as the Hebrew noun, producing contextually impossible readings. The LLM detected 36 such errors by recognizing the named entities from its training data and flagging the semantic incongruity.

In a complementary pattern, the homograph די (/dj/; in Hebrew usually “enough”/“quite”) encompasses foreign names where the correct vowel depends on the actual pronunciation of the person’s name: Ruby *Dee* and Chuck *D* require די (/di/, with *hiriq*), while Daniel *Day*-Lewis requires די (/dej/, with *tsere*). Across two files, 29 such errors were confirmed, each requiring the model to identify the specific individual and the phonetics of the individual’s name.

#### 6.4 The אס Case Study: Asymmetric Precision

A major source of false positives is the consonantal string אס (/ʔm/), corresponding to אס (/ʔem/, “mother”) or אס (/ʔim/, “if”).

In our analysis, the two אס files were outliers with notably low precision of the LLM runs (36.5% and 33.1% respectively), accounting for the bulk of spurious flags. The main issue is the Hebrew transliteration of the Latin letter *M*. The LLM correctly identified occurrences of “M” as semantically distinct from the Hebrew word for “mother,” but suggested diacritizing it with *segol* (אס) rather than *tsere* (אס). The two vowels are phonetically identical in Modern Hebrew; the distinction is purely orthographic and governed by the Academy of Hebrew Language, which prescribes *tsere*. We note that the morphology list supplied to the model contained only the canonical Hebrew readings (“mother”, “if”) and not the Latin letter-name use, leaving the model without an in-prompt anchor for the prescribed form. While such an entry could in principle be added, our ablation (§5.2) shows that morphological context yields only marginal gains, and the deeper issue, such as disagreement among prescriptive norms regarding loan words, would persist regardless.

This case illustrates a fundamental challenge in auditing evaluation: the “correctness” of a diacritization depends on the prescriptive norm adopted. Under a strict Academy norm, the LLM’s suggestion is wrong (FP), but the *segol* may be considered acceptable by other audiences. Despite the low precision, the model’s ability to distinguish the foreign-letter meaning from the Hebrew word demonstrates genuine semantic understanding, even though its specific vowel suggestion does not match the prescribed form.

## 6.5 Summary of Detected Errors

In the focused subset (15 files, 7,241 sentences), the 271 confirmed errors correspond to an error rate of 3.7%. The union of all 50 focused runs detected 266 of these (98.2%), with the 5 remaining errors discussed in §6.2. Two full-corpus runs (Gemini 3 Pro, B200–m–d, and B500, +m+d) covering all 332 files in the Challenge Set flagged a combined 4,731 instances outside the current GS. Our human domain expert has adjudicated each of these instances, thus producing a new and improved version of the EACL 2023 Challenge Set. We hereby release this new version of the Challenge Set, reflecting all errors confirmed to date.

## 6.6 Workflow Comparison: LLM-Aided vs. Manual Auditing

To validate our Gold Standard and quantify the practical value of LLM-based auditing, we compare two annotation-review workflows on the same 15 focused files (7,241 sentences, 271 confirmed errors):

- **Manual auditing:** Two independent native Hebrew speakers with graduate-level linguistics training (Expert 1 and Expert 2) each reviewed all 7,241 sentences sequentially (XLSX files, one row per sentence, target diacritization highlighted), without access to model output or each other’s annotations. Each review required approximately 30 hours over one week; both experts were compensated at standard rates.
- **LLM-aided auditing:** Gemini 3 Pro (batch size 200, without morphological or dictionary data, averaged over 2 runs) processed the same sentences via API. Each run completed in ~75 minutes of wall-clock time at a cost of ~\$7; the human reviewer then verified only the flagged findings (~225 items per run).

Expert 1 achieved 94.9% precision and 62.4% recall; Expert 2 achieved 95.9% precision but only 42.8% recall. Both maintained high precision (>94%), yet their recall differed by 20 percentage points, highlighting the inherent inconsistency of manual error detection in sparse-target settings. In contrast, the LLM-aided workflow achieved 83.6% recall with 99.1% precision (Figure 2). The recall gaps of 21.2 pp (vs. Expert 1) and 40.8 pp (vs. Expert 2) are both statistically significant (McNemar’s test,  $p < 0.001$ ).

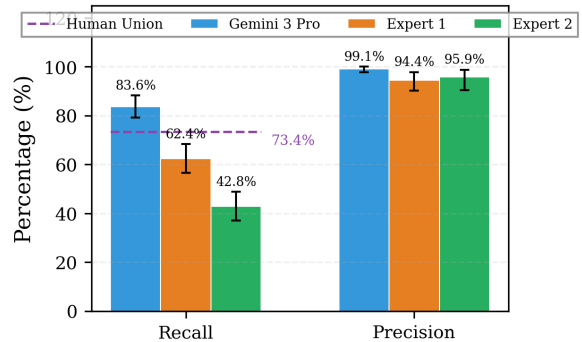


Figure 2: Workflow comparison: LLM-aided auditing (Gemini 3 Pro, batch size 200, no auxiliary context, averaged over 2 runs) vs. two independent human experts. Recall and precision are measured against our 271-error gold standard. The dotted line marks the union recall of both experts combined (73.4%), which still falls short of the LLM’s 83.6%. Error bars indicate 95% confidence intervals.

Even the **union** of both experts’ findings reaches only 73.4% recall (199/271), still 10.2 pp below the LLM’s 83.6%. This result demonstrates that the recall gap is not attributable to one annotator’s idiosyncratic weakness: combining two qualified reviewers still fails to match the automated pipeline.

The gap (Figure 2) does not imply that the LLM “understands” Hebrew better than the experts; rather, it reflects the inherent difficulty of sustaining attention across thousands of mostly-correct sentences. At the observed error density of 3.7% (271/7,241), even a diligent reviewer encounters one genuine error roughly every 27 sentences. This is a sparse-target search task where human fatigue is a well-documented limitation (Swayamdipta et al., 2020).

The LLM-aided workflow shifts the reviewer’s task from *open-ended search* (scanning 7,241 sentences for anomalies) to *targeted verification* (adjudicating ~230 flagged items), reducing review effort by over 95% while improving recall. This supports the practical recommendation that LLM-based auditing is most valuable not as a replacement for human expertise, but as a **triage layer** that focuses human attention where it is most needed.

## 6.7 Inter-Annotator Reliability

On the binary sentence-level task (error or not error, Cohen’s  $\kappa = 0.571$  (“moderate”)), but this low value is an artifact of extreme class imbalance: with only 3.7% error prevalence, raw agreement is 98.3%, a well-known prevalence paradox where  $\kappa$  underestimates true agreement (Feinstein and Cicchetti,

1990; Cicchetti and Feinstein, 1990).

Of the 271 confirmed errors, only 86 (31.7%) were found by both experts; 83 were found only by Expert 1, 30 only by Expert 2, and 72 (26.6%) by neither. This low overlap reinforces that human error detection in sparse-target settings depends less on linguistic competence than on sustained vigilance. In contrast, repeated LLM runs exhibit much lower variance: two Gemini 3 Pro runs at identical settings yielded 225 and 228 TP respectively (recall spread: 1.1 pp vs. 20 pp for humans).

## 7 Post-Hoc Quality Filters

Given the substantial gap in false-positive rates across models (Table 1), we investigated several post-hoc filtering strategies. Since our task is recall-centric (surfacing as many real errors as possible for human review), filtering is primarily useful when precision is low enough to impose a significant manual review burden.

### 7.1 Finding-Level Filtering

We investigated two post-hoc filters for reducing false positives. **Confidence filtering**, i.e. retaining only High-confidence findings, has negligible effect on Gemini 3 Pro (which already labels nearly all findings as High) but raises Flash precision from 39.7% to 64.0% at a steep recall cost (−15.5 pp). **Trivially invalid findings**, where the suggested correction is identical to the original diacritization (after Unicode normalization) or empty, account for 25.3% of Flash output but under 1% of Gemini Pro output; filtering them improves Flash precision with near-zero recall loss. For the best-performing model, neither filter has a meaningful effect, as its precision already exceeds 99%.

### 7.2 Multi-Run Union: Trading Precision for Recall

Since our workflow prioritizes recall, surfacing errors for human review, we analyze how combining the output of multiple runs affects coverage. Taking the **union** of findings from two or more runs increases recall, but the precision cost depends strongly on which runs are combined. Same-model Gemini Pro unions preserve high precision, while cross-model and all-run unions surface many more false positives.

Table 5 and Figure 3 present practical multi-run strategies that a user could employ without prior knowledge of which specific runs are optimal. A

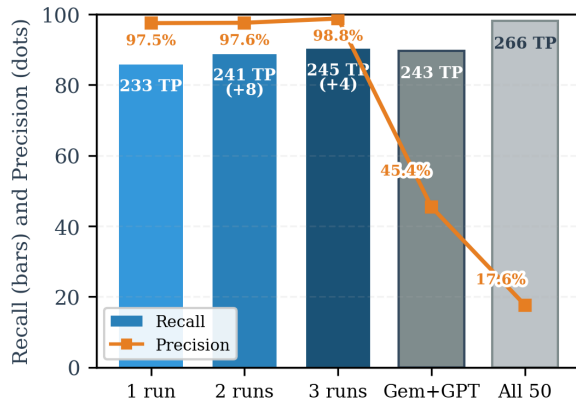


Figure 3: Diminishing returns of multi-run union. Bars show recall with True Positive counts; the orange line shows precision.

single Gemini 3 Pro run achieves up to 86.0% recall (batch size 100, with morphology data but without dictionary data). Adding a second run, whether a repeat of the same configuration, a different configuration of the same model, or a different model entirely, raises recall to ~89%. A third run of the same model in a different configuration reaches 90.4% recall while maintaining 98.8% precision. With batch size 50, all three diversification strategies yield similar marginal gains (mean +9–11 TP across Gemini Pro B50 pairs), suggesting that the stochastic variation between runs is the primary driver of union gains, not the specific form of diversity.

Strategy	TP	Recall	Prec.	F1
Single best run (Gem. Pro B100)	233	86.0%	97.5%	91.4%
2× same config (repeat)	241	88.9%	97.6%	93.1%
2× diff. config (same model)	241	88.9%	97.6%	93.1%
Gem. Pro + GPT-5.2 (B50)	243	89.7%	45.4%	60.3%
3× Gem. Pro (mixed B & config)	245	90.4%	98.8%	94.4%
All 50 runs (theoretical ceiling)	266	98.2%	17.6%	29.9%

Table 5: Recall, precision, and F1 from union of multiple runs (out of 271 Gold Standard errors). Same-model Gemini Pro unions raise recall to ~89–90% while maintaining high precision; cross-model and all-run unions recover additional errors but introduce many more false positives. Five errors (including 3 dormant) were never detected by any of 50 runs. Gem. Pro = Gemini Pro; B = batch size.

The cross-model union with GPT-5.2 adds coverage beyond Gemini Pro, detecting 13 errors missed by the Gemini Pro B50 run and reaching 89.7% recall, but at a substantial precision cost (45.4%). The theoretical ceiling across all 50 runs remains 98.2%; however, its precision is only 17.6%, so

this setting should be interpreted as a recall upper bound rather than a practical review strategy. The 5 remaining errors consist of 2 ambiguous cases discussed in §6 plus 3 dormant errors that eluded all models and both human experts alike.

For practical deployment, we recommend **two to three runs of the same model** in different configurations as a cost-effective strategy: this raises recall to  $\sim 89\text{--}90\%$  while maintaining precision above  $97\%$ .

## 8 Future Directions

The auditing pipeline presented in this work operates on a focused confusion-set structure: each sentence contains a single pre-identified homograph that the model is asked to evaluate. This setup maximizes per-query precision and allows controlled evaluation, but it presupposes that the set of candidate homographs is known in advance. Three natural extensions suggest themselves: **running-text auditing**, in which errors are flagged in arbitrary text without pre-specified target words; **homograph diacritization**, in which LLMs are used to predict the correct diacritized form of a homograph in context rather than merely validating existing labels; and **full running-text diacritization**, applying LLMs to the general task of adding *nikud* to undiacritized Hebrew text, which subsumes both of the preceding extensions and is likely the hardest of the three. Preliminary single-run probes for all three extensions are detailed in Appendix H; they suggest that frontier LLMs retain high precision on running text but pay a substantial recall cost without the focusing effect of specific words, that diacritization-from-scratch is already strong on easier homograph cases, and that on full running-text diacritization a frontier LLM trails the specialist *Dicta* auto-diacritizer overall ( $92.0\%$  vs.  $95.4\%$ ) yet edges it at homograph positions. Full characterization of these extensions is left for future work.

## 9 Conclusion

We presented a human–AI auditing workflow for identifying Hebrew *nikud* errors, achieving high precision and recall with frontier LLMs. Our findings demonstrate that performance is primarily sensitive to batch size, whereas auxiliary morphological and dictionary context does not yield consistent improvements under zero-shot conditions, suggesting that frontier LLMs possess sufficient internal linguistic knowledge for this task. A workflow com-

parison with two independent human experts shows that manual recall varies widely ( $42.8\%$ – $62.4\%$ ) and that even the union of both experts’ findings ( $73.4\%$ ) falls short of a single LLM run ( $83.6\%$ ). In contrast, our LLM-aided auditing pipeline reduces review effort by over  $95\%$ , shifting the human role from open-ended search to targeted verification, improving both recall and cost. Given the manageable inference costs and high effectiveness, this pipeline is economically viable for large-scale dataset maintenance. We release our Gold Standard dataset, including per-error difficulty annotations derived from cross-system detection rates, as well as the full corrected EACL 2023 Challenge Set to support further work within the Hebrew NLP community.

## Limitations

While our auditing workflow demonstrates high precision and recall, several factors constrain the interpretation of these metrics.

First, the definition of a diacritization error is subject to **prescriptive linguistic norms**. As shown in the case of the consonantal string  $\text{מס}$  (serving as the letter *M*), a single decision, such as whether a loanword requires *tzere* or *segol*, can shift the measured precision of specific files by over 50 percentage points. Our evaluation assumes the conventions of the source dataset, yet we acknowledge that alternative normative decisions would lead to different performance profiles. This is further compounded by **orthographic noise** (typos) in the source corpus; both human and AI auditors often identify inconsistencies that stem from spelling errors rather than diacritization issues, which technically fall outside the scope of our task but impact the measured precision.

The second factor is the **dormant-error problem**. Although we tried to mitigate this through exhaustive manual review (§4.4), and through the validation of the Gold Standard via a second independent annotator (§6.7), the final decisions remain those of a single adjudicator. The fact that Expert 2’s 5 disagreements were all resolved in favor of the existing ground truth provides evidence of adjudication robustness, but errors subtle enough to elude both annotators and all models may remain undiscovered; the  $0.09\%$  dormant rate we report is therefore itself an upper bound.

A related concern is **Gold Standard circularity**: since model runs contributed to the initial candidate pool used to construct the Gold Standard

(§4.4), errors that no model detects are structurally underrepresented in the denominator, potentially inflating measured recall. We mitigate this in three ways: (1) two independent human experts flagged errors without access to model output, contributing candidates that models missed; (2) exhaustive manual review of all 6,378 non-flagged sentences uncovered 6 additional dormant errors; and (3) the Gold Standard explicitly includes 5 errors that no model detected across all 50 runs, confirming that the denominator is not limited to model-detectable cases. Nevertheless, we acknowledge that if a class of errors exists that is systematically invisible to both humans and models, it would remain absent from the Gold Standard, and true recall would be lower than reported.

Finally, our reliance on **proprietary frontier LLMs** poses challenges for long-term reproducibility. API updates and shifting cost structures may affect the consistency of auditing results over time. In future work we plan to explore the viability of fine-tuned open-source models to ensure transparent and repeatable dataset maintenance.

## 10 Acknowledgments

This work has been funded by the Israel Science Foundation (grant No. 2617/22) and by the European Union (ERC, MiDRASH, Project No. 101071829; principal investigators: Avi Shmidman, Bar-Ilan University; Daniel Stökl, EPHE-PSL; Nachum Dershowitz, Tel Aviv University; and Judith Olszowy-Schlanger, EPHE-PSL), for which we are grateful. Views and opinions expressed are, however, those of the authors only and do not necessarily reflect those of the European Union or the European Research Council Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

## References

- Yaacov Choueka and Yoni Neeman. 1995. Nakdan-text, (an in-context text-vocalizer for modern hebrew). In *BISFAI-95, The Fifth Bar Ilan Symposium for Artificial Intelligence*.
- Domenic V. Cicchetti and Alvan R. Feinstein. 1990. [High agreement but low kappa: II. Resolving the paradoxes](#). *Journal of Clinical Epidemiology*, 43(6):551–558.
- Alvan R. Feinstein and Domenic V. Cicchetti. 1990. [High agreement but low kappa: I. The problems of two paradoxes](#). *Journal of Clinical Epidemiology*, 43(6):543–549.
- Ya’akov Gal. 2002. An hmm approach to vowel restoration in arabic and hebrew. In *Proceedings of the ACL-02 workshop on Computational approaches to semitic languages*, pages 1–7. Association for Computational Linguistics.
- Elazar Gershuni and Yuval Pinter. 2022. [Restoring Hebrew diacritics without a dictionary](#). In *Findings of the Association for Computational Linguistics: NAACL 2022*, pages 1010–1018, Seattle, United States. Association for Computational Linguistics.
- Google. 2026. [Gemini 3 — API documentation](#). Accessed: 2026-02-16.
- Curtis G. Northcutt, Lu Jiang, and Isaac L. Chuang. 2021. [Confident learning: Estimating uncertainty in dataset labels](#). *Journal of Artificial Intelligence Research*, 70:1373–1411.
- Avi Shmidman, Cheyn Shmuel Shmidman, Dan Bareket, Moshe Koppel, and Reut Tsarfaty. 2023. [Do pre-trained contextual language models distinguish between Hebrew homograph analyses?](#) In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 831–843, Dubrovnik, Croatia. Association for Computational Linguistics.
- Avi Shmidman, Shaltiel Shmidman, Amir Eyal, and Moshe Koppel. 2020. [Nakdan: Professional Hebrew diacritizer](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 197–203. Association for Computational Linguistics.
- Swabha Swayamdipta, Roy Schwartz, Nicholas Lourie, Yizhong Wang, Hannaneh Hajishirzi, Noah A. Smith, and Yejin Choi. 2020. [Dataset cartography: Mapping and diagnosing datasets with training dynamics](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 9275–9293, Online. Association for Computational Linguistics.

## A Full-Sentence Examples from the Dataset

This appendix presents four representative sentences from the focused evaluation set, along with English translations.

### Example 1: Part-of-Speech (Noun → Verb).

אמא שלי «אָהבָה» רקגבר אחד מכל מי שהיה לי והוא כרגע בן 24 בדרך לחתונה והיא עדיין מחכה שנחזור.

*Original:* אָהבָה ('ahava, noun “love”)  
*Correction:* אָהבָה ('ahava, verb 3FS.PST “she loved”)

*Translation:* “My mother **loved** only one man out of all those I had, and he is currently 24, on his way to a wedding, and she is still waiting for us to come back.”

*Notes:* The two forms are phonetically identical in Modern Hebrew; only the contextual reading (subject-verb-object) reveals the verbal interpretation.

### Example 2: Syntax (Adjective → Construct Noun).

הוא מוכר למשטרה מעבירות «אלימות» הקשורות לפעילותו בימין הקיצוני.

*Original:* אֲלִימוֹת ('alimot, adjective F.PL “violent”)  
*Correction:* אֲלִימוֹת ('alimut, noun F.SG “violence”)  
*Translation:* “He is known to the police for **violence** offenses connected to his activity on the far right.”

*Notes:* The standard legal term is the construct compound “offenses of violence” (N+N), not the attributive “violent offenses” (N+Adj).

### Example 3: Inflection (Plural Construct → Singular Absolute).

אבל אחדות נגד אויב משותף זה «תנאים» בסיסי לחברה, אבל זה לא מספיק...

*Original:* תְּנָאִי (tna'ei, plural construct “conditions of”)

*Correction:* תְּנָאִי (tnai, singular absolute “a condition”)

*Translation:* “But unity against a common enemy is a basic **condition** for society, but it is not enough...”

*Notes:* The agreement cues, singular copula זה and singular adjective בסיסי, require a singular noun.

## Example 4: Morphology (Construct → Possessive Suffix).

ובך, גם אם ניסיתי שלא לערב את עבודתי «בְּחַיִּי» האישיים, נשזר סיפור חוויותי המקצועיות באופן טבעי בתולדות חיי.

*Original:* בְּחַיִּי (be-hayyei, prep. + construct plural “in the lives of”)

*Correction:* בְּחַיִּי (be-hayyai, prep. + plural with 1SG possessive “in my life”)

*Translation:* “Even though I tried not to involve my work **in my** personal life, the story of my professional experiences became naturally woven into my life history.”

## B System Prompt

The models were prompted using a linguist-persona template designed to prioritize recall over precision (Appendix D). Output was constrained to a structured JSON schema (Appendix E).

## C API Request Structure

This appendix presents a concise overarching view of the request structure; the full text of the prompt is given below in Appendix D.

```
— System Prompt (Appendix D) —
You are an expert linguist specializing in Hebrew grammar, morphology, syntax, and particularly nikud...

— Morphological Data (optional) —
אָהבָה: Lemma=אָהַב, <Verb Bareinfinitive Paal>
אָהבָה: Lemma=אָהַב, <Verb Fem Sg P3 Past Paal>
אָהבָה: Lemma=אָהַב, <Noun Fem Sg Absolute>
...(31 diacritized forms total)

— Dictionary Definitions (optional) —
אָהַב
רגש עמוק של משיכה לאדם או לדבר
אָהַב (Qal)
נמשך ברגש עמוק לאדם או לדבר
...(3 entries total)

— Sentences (batch of 50–500) —
S1: ... הוא חוצה את ... שאשתו «אָהַבָה» במיוחד.
S2: היא «אָהַבָה» את התולה מריה.
...
```

Figure 4: Structure of a single API request, illustrated with the homograph אָהַב. Each batch includes the system prompt, optional morphological and dictionary context, and 50–500 sentences. The target word is marked with guillemets.

## D Prompt Text

The following is the complete system prompt sent to each model. The prompt was authored in Markdown; formatting conventions (***\*\*bold\*\****→**bold**, ***\*italic\****→*italic*, ``code``→code) are rendered here as their visual equivalents.

You are an expert linguist specializing in Hebrew grammar, morphology, syntax, and particularly *nikud* (vocalization). Your task is to act as a quality assurance agent for Hebrew sentences where one specific word form has been vocalized.

### Task Context:

You will receive a list of sentences. In **all** these sentences, the same string of letters has been vocalized in the **exact same way** (marked with `<<...>>`). Your goal is to validate if this specific vocalization is contextually correct for each sentence. Since Hebrew homographs (words with same letters but different vowels/meanings) are common, the provided vocalization might be correct for some sentences but incorrect for others (where a different vocalization of the same letters is required).

### Key Objective:

Achieve high recall of actual *nikud* errors. It is acceptable if up to 20% of your reported errors are false positives. If you are uncertain but suspect an error based on context, report it with Medium or Low confidence.

### Input Structure:

1. **Morphological Data:** This data lists unvocalized words, their possible *nikud* forms, and corresponding lemma/grammatical information (e.g., `<Lemma=...>`, `<Gram=...>`). This data is extensive but not necessarily complete (especially for loanwords or rare forms) and may, very rarely, contain errors. Use this list as a reference for possible alternative vocalizations if the provided one seems wrong.
2. **Sentences:** A list of sentences (e.g., S1, S2, ...), each with the target word marked by `<<...>>`.

### Processing Instructions for Each Sentence:

1. Identify the unvocalized letters of the word within `<<...>>`.
2. Carefully analyze the full sentence context surrounding the marked word. Pay attention to syntax, semantics, and grammatical agreement.
3. **Validate the Provided Nikud:**
  - Does the specific form provided in `<<...>>` fit the sentence grammatically and semantically?
  - **If YES:** Move to the next sentence.
  - **If NO:** This is an error. You must determine the correct vocalization.
    - Consult the **Morphological Data**: Does another form in the list fit the context perfectly?
    - If the correct form is not in the data (rare), derive it based on your linguistic knowledge.

### 4. Contextual Validation – Specific Checks:

- **Part of Speech & Tense:** This is the most common error in this task. For example, if the input is `<<אֵל>>` (Present/Participle), but the sentence says `“...מִחֵר...”`, the context requires Future tense (`<<אֵל>>` or `<<אֵלֶיךָ>>`).
  - **Gender and Number Agreement:** Does the *nikud* agree with other elements in the sentence?
  - **Construct State (*Smichut*) vs. Absolute State:** Is the word part of a *smichut*? Is the vowel pattern correct for that state?
  - **Definiteness with Prefixes (ב, ג, ל):** Check if the vowel under the prefix indicates the correct definiteness (e.g., `<<בְּסֵפֶר>>` vs `<<בְּסֵפֶרֶת>>`).
  - **Loanwords and Proper Nouns:** Use world knowledge for names or foreign terms not fully covered in morphology.
  - **Original orthography:** If the sentence includes `[[[original: ...]]]`, consider the spelling changes when determining the correct vocalization.
5. **What NOT to Flag as Errors:**
    - **Shva vs. Hataf:** Do NOT report differences like `<<מִעֵן>>` vs `<<מִעֵןֶת>>` unless strictly necessary.
    - **Minor Stylistic Variants:** Acceptable variants are not errors.
    - **The Unvocalized Word Itself:** Assume the base letters are correct.
  6. **Reporting Errors:**
    - If the provided *nikud* is contextually incorrect:
      - Populate the JSON object.
      - In `suggested_correct_nikud_form`, provide the correct vocalization (preferably from the Morphological Data if a match exists).
      - In `reason_for_error`, explain the mismatch (e.g., “Context implies Future tense 1st person, but word is vocalized as Present tense”).

### Output Format:

You must output a JSON object strictly conforming to the following schema. The top-level object should have a single key: “errors”, which is an array of error objects. If no errors are found, this array should be empty.

## E Output JSON Schema

Each model was instructed to return a structured JSON object conforming to the following schema.

```
{
  "type": "object",
  "properties": {
    "errors": {
      "type": "array",
      "items": {
        "type": "object",
        "properties": {
          "sentence_id": {
            "type": "string",
            "description": "Sentence ID (e.g., S1, S2)"
          },
          "full_sentence": {
            "type": "string",

```

```

    "description": "Complete Hebrew
      sentence as provided"
  },
  "marked_text_unvocalized": {
    "type": "string",
    "description": "Unvocalized letters
      of the marked word"
  },
  "provided_nikud_form": {
    "type": "string",
    "description": "Vocalized word as
      it appeared in input"
  },
  "suggested_correct_nikud_form": {
    "type": "string",
    "description": "Suggested correct
      vocalized form"
  },
  "reason_for_error": {
    "type": "string",
    "description": "Explanation of why
      the nikud is incorrect"
  },
  "confidence_in_error": {
    "type": "string",
    "enum": ["High", "Medium", "Low"]
  },
  "morphological_data_feedback": {
    "type": "string",
    "description": "Optional feedback
      on morphological data gaps"
  }
}
},
"required": [
  "sentence_id",
  "full_sentence",
  "marked_text_unvocalized",
  "provided_nikud_form",
  "reason_for_error",
  "confidence_in_error"
]
}
},
"required": ["errors"]
}

```

## F Recall on Injected Errors

To validate recall measurement under controlled conditions, we constructed contaminated files by injecting synthetic errors into cleaned homograph files. Starting from 5 files spanning 3 homograph groups (2,360 sentences total), we swapped 10% of homograph tokens with a competing diacritized alternative selected at random, producing 308 injected errors (Table 7).

Model	Prec.	Rec.
Gemini 3 Pro	100.0	78.9–100.0
GPT-5.2	99.3–99.7	94.5–98.4
Gemini 3 Flash	96.8–97.7	82.5–96.8

Table 6: Contaminated-set results (5 files, 308 injected errors).

## G Model Identifiers and Configuration

Table 8 lists the API identifiers for each model evaluated in our study. All models were accessed via the OpenRouter API between January 25 and February 16, 2026; exact version snapshots are not avail-

Homograph file	Sentences	Injected
הַדָּוָד_1	489	46
הַדָּוָד_2	514	64
דָּוָד_1	497	44
דָּוָד_2	508	59
הַדָּוָד	352	95
Total	2,360	308

Table 7: Contaminated-set composition: sentences per file and injected errors (10% swap rate). Suffixes \_1/\_2 denote sub-files partitioned from the original dataset for batch processing.

able through the API. The prompt was not individually optimized per model architecture; performance differences may therefore partly reflect prompt compatibility rather than intrinsic model capability.

Model	API Identifier
Gemini 3 Pro	gemini-3.0-pro-preview
Gemini 3 Flash	gemini-3.0-flash-preview
GPT-5.2	gpt-5.2
Claude Opus 4.5/4.6	claude-opus-4-*

Table 8: Model API identifiers. All models accessed via OpenRouter; exact version snapshots are not available through the API.

**Token Usage.** Table 9 reports the total input and output token counts for each model–batch-size combination on the focused evaluation set. Input counts include the system prompt repeated once per batch; larger batch sizes amortize this overhead across more sentences, resulting in lower total input tokens. These counts allow readers to recompute costs under updated pricing.

Model	Batch	Input	Output	Total
Gemini 3 Pro	50	663,600	945,103	1,608,703 <sup>†</sup>
	200	526,530	491,989	1,018,520 <sup>†</sup>
	500	485,055	283,539	768,594 <sup>†</sup>
GPT-5.2	50	608,436	326,395	934,831
	200	451,248	130,335	581,583
Claude Opus 4	50	812,421	60,935	873,356
	200	693,808	68,053	761,861
	500	680,159	24,277	704,436
Gemini 3 Flash	50	616,893	172,923	789,816
	200	492,533	65,929	558,462
	500	478,431	27,571	506,002

Table 9: Token usage per run on the focused evaluation set (7,241 sentences). Input and output counts are totals across all API calls in a single run. <sup>†</sup>Averaged over multiple configurations. These counts allow cost recomputation under updated pricing schedules.

## H Preliminary Probes

This appendix reports three single-run probe studies that empirically anchor extensions discussed in the Future Directions section. We report these as single-run probe numbers; full characterization is left for future work. The main experiments in this paper were conducted in January–February 2026 with the `gemini-3.0-pro-preview` snapshot, which has since been retired; the probes below were run in May 2026 using its closest currently-available successor, `gemini-3.1-pro-preview`.

### H.1 Running-Text Auditing

We ran a single-run probe on 77 sentences drawn from the focused corpus (20 each from the Easy/Medium/Hard error tiers and 20 randomly-sampled Correct sentences, with three removed for duplication or off-domain content), stripped of guillemets and diacritic marks, and re-diacritized by *Dicta Nakdan* (Shmidman et al., 2020) as an off-the-shelf auto-diacritized baseline. The first author manually corrected Dicta Nakdan’s output to produce a blind reference; one additional Gemini-flagged error missed in the blind pass was promoted post-hoc, giving 71 blind and 72 adjudicated gold errors. Dicta Nakdan resolved most originally-marked homographs without guillemets (only  $\sim 10$  of the original homograph positions remained wrong post-diacritization), but introduced 62 new errors elsewhere in the text. Using `gemini-3.1-pro-preview` with `morph` and `dictionary` headers disabled, single-run inference yields whole-text  $P = 100\%$ ,  $R = 41.7\%$ ,  $F_1 = 0.59$  over 30 flagged errors and 72 gold; correction-exact-match among true positives is 76.7%. Stratified, recall is 40% (4/10) at original homograph positions, a substantial drop from the focused pipeline’s 83.6%, and 41.9% (26/62) on Nakdan-introduced errors elsewhere, suggesting that the recall cost of removing guillemets is comparable in both regions of the text. Failure modes cluster around foreign-name transliteration, definiteness on  $\text{כ/ל}$  prefixes, and recurring focused-corpus homographs flagged in some sentences but missed at other running-text occurrences ( $\text{אֲנִי, אֲנִי}$ ). Two caveats temper this picture: the gold is human-adjudicated against model output and is not an exhaustive independent benchmark, and the homograph-position subset ( $n = 10$ ) and per-tier counts are too small for reliable per-tier conclusions.

### H.2 Homograph Diacritization

We ran a single-run probe on 542 sentences (all 271 confirmed errors plus a per-file matched Correct sample, randomly ordered across mixed-confusion-set batches) using `gemini-3.1-pro-preview` at the same B200 batching budget. Without `morph` context, the model produces the gold diacritized form on 88.6% of the 271 error positions and preserves the published form on 98.5% of matched correct positions, with 100% consonant preservation and no missing predictions. Two caveats temper this picture: a batch-level union `morph` header, which substitutes for the per-file header used in our auditing setup (§5.2) given the mixed batches, in fact *hurts* error-correction by 10.7 pp, with the regression concentrated in prefix-letter homographs (e.g.,  $\text{א}$  vs.  $\text{א}$  in  $\text{בְּהַמְשֵׁךְ, בְּהַחַי}$ ) where the union header appears to bias the model away from definite-article-bearing forms; and the Hard difficulty band stalls at 51.2% in both configurations, marking cases where context-dependent decisions remain robustly hard.

### H.3 Full Running-Text Diacritization

We ran a single-run probe on the same 80 sentences used in §H.1, asking `gemini-3.1-pro-preview` (without `morph` or `dictionary` headers) to diacritize each token from scratch as plene undiacritized Hebrew text, and compared the result token-by-token to the human-adjudicated gold and to *Dicta Nakdan* as an off-the-shelf auto-diacritized baseline. Gemini reaches 92.0% across  $n = 1,551$  active tokens versus Dicta Nakdan’s 95.4%; at the 76 original homograph positions Gemini’s 88.2% slightly edges Dicta Nakdan’s 86.8%, and on the Hard tier Gemini retains a substantial lead (81.8% vs. 68.2%,  $n = 22$ ), the same context-dependent pattern as in §H.2. Two caveats temper this picture: the 80-sentence, 1,551-token sample is small and shares provenance with §H.1, so token-level scores are not an independent benchmark of running-text *nikkud*; and the Hard-tier subset ( $n = 22$ ) is suggestive rather than conclusive. A larger, independently-curated full-text gold, evaluated blind to both systems, would be needed for a robust Nakdan-vs.-LLM verdict.

# Beyond Annotator Disagreement: Guideline-Induced Errors in Arabic Hate Speech Annotation

Wajdi Zaghouni

Northwestern University Qatar  
wajdi.zaghouni@northwestern.edu

## Abstract

Annotation errors in hate speech corpora are often attributed to annotator disagreement or bias. This paper argues that a substantial and underexamined class of errors originates upstream, from structural weaknesses in annotation guidelines themselves. When guidelines fail to encode the linguistic and cultural properties of the target discourse, they make certain errors structurally inevitable regardless of annotator quality. Focusing on Arabic social media discourse, a challenging setting due to its dialect continuum, culturally embedded insult conventions, sarcasm-heavy pragmatics, and complex religious rhetoric, we identify three mechanisms through which guideline design produces systematic annotation errors: cultural misclassification, when culturally specific hostile expressions fall outside annotation categories; dialectal ambiguity, when lexical meanings shift across regional varieties; and annotation projection, when frameworks developed for English moderation are applied to Arabic without adequate adaptation. Using six illustrative case studies with attested Arabic examples, we show how these mechanisms produce recurrent misannotations in existing datasets. We propose a taxonomy of five guideline-induced error types, an explicit mapping from mechanisms to error types, and a practical four-stage diagnostic framework for dataset builders.

## 1 Introduction

Annotation errors are a persistent challenge in NLP dataset construction, and their sources are not always well understood.

In the hate speech detection literature, annotation inconsistencies are most frequently attributed to individual annotators: their disagreements, biases, or insufficient training (Waseem and Hovy, 2016; Davidson et al., 2017). Inter-annotator agreement (IAA) metrics such as Cohen’s  $\kappa$  and Krippendorff’s  $\alpha$  are routinely used to assess reliabil-

ity (Artstein and Poesio, 2008). While such metrics are informative, recent work has emphasized that they conflate two distinct phenomena. Pavlick and Kwiatkowski (2019) demonstrate that annotator disagreement often reflects genuine linguistic uncertainty rather than annotation failure. Nie et al. (2020) further show through large-scale collection of multiple annotations per item that meaningful distributions of human opinion persist even on widely used NLI evaluation sets, and Jiang et al. (2023) provide ecologically valid explanations for why such variation arises. Weber-Genzel et al. (2024) introduce a rigorous methodology for separating annotation errors from human label variation, establishing that the two require different remediation strategies. Klie et al. (2023) survey annotation error detection methods across eighteen approaches and nine datasets, showing that many widely used corpora contain substantial numbers of objectively incorrect labels.

This paper contributes to that line of inquiry by examining a neglected *upstream* source of annotation error: the guidelines themselves. We argue that when annotation guidelines fail to account for the linguistic and cultural properties of the target discourse, they systematically produce errors regardless of annotator quality. These are not errors caused by annotator sloppiness or disagreement; they are errors *structurally forced* by the design of the annotation scheme, inscribed before a single label is assigned.

We develop this argument in the context of Arabic social media discourse, which is a particularly vulnerable setting for guideline-induced errors. Four properties compound the risk: (1) a *dialect continuum* in which the same word carries sharply different meanings across Gulf, Levantine, Egyptian, and Maghrebi varieties (Habash, 2010; Darwish et al., 2017); (2) *culturally embedded hostility* targeting lineage, tribal affiliation, and sectarian identity with no equivalent in English-derived

categories (Mubarak et al., 2017; Albadi et al., 2018); (3) *sarcasm-heavy pragmatics* in which conventional praise encodes contempt; and (4) *religious rhetoric* in which intra-sectarian delegitimization is misclassified by generic religion categories. Each property corresponds to a mechanism through which guideline inadequacy becomes annotation error.

Two clarifications are warranted before proceeding. First, while sarcasm, irony, and culturally specific insults exist in many languages and cultures, what matters here is not their existence in isolation but the specific way they interact with Arabic dialect distribution and with the categorical schemata of dominant English language guidelines. Second, dialectal variation is a general challenge for Arabic NLP, but in the context of hate speech annotation it has a specific consequence we make precise below: a single token can be hostile in one dialect and benign in another, so a guideline that does not stratify by dialect produces systematic disagreement that is correlated with annotator background rather than with the text under annotation.

We identify three mechanisms through which annotation guidelines produce systematic errors in this setting:

1. **Cultural misclassification:** culturally specific hostile expressions are absent from annotation categories, leading to inconsistent labeling across annotators who each follow the guidelines as written.
2. **Dialectal ambiguity:** lexical items whose meaning varies across Arabic dialects are assigned labels based on only one of their possible interpretations, producing systematic disagreement correlated with annotator dialect background.
3. **Annotation projection:** annotation schemes designed for English social categories are transferred to Arabic without adaptation, introducing categorical mismatches that no annotator can resolve within the given framework.

Through linguistic case studies with attested Arabic examples, we illustrate each mechanism and show how it generates recurrent annotation errors. We then propose a five-type taxonomy of guideline-induced errors, an explicit mapping from the three mechanisms to the five error types, and a four-stage diagnostic framework.

## 2 Background

### 2.1 Annotation Errors and Human Label Variation

Artstein and Poesio (2008) provide a foundational survey of inter-coder agreement measures and the conditions under which each is appropriate. High IAA has been treated as a proxy for annotation quality, but this equivalence has been repeatedly challenged.

Plank et al. (2014) argue that disagreements in NLP annotation are often linguistically meaningful rather than spurious noise. Pavlick and Kwiatkowski (2019) show that disagreement in NLI persists even with additional context, suggesting real linguistic uncertainty rather than annotator error. Nie et al. (2020) collect ChaosNLI and find that genuine human disagreement is widespread and that current models fail to capture the resulting label distributions. Jiang et al. (2023) elicit ecologically valid explanations from annotators to identify the linguistic and pragmatic sources of label variation, and Giulianelli et al. (2023) broaden the picture to NLG, connecting human production variability to aleatoric uncertainty. Taken together, this body of work establishes that human label variation (HLV) is a pervasive and informative signal, not a defect to be eliminated. Swayamdipta et al. (2020) introduce dataset cartography, a model-based approach that reveals many annotated datasets contain a substantial proportion of ambiguous or mislabeled examples.

Most directly relevant to our work, Klie et al. (2023) systematically analyze annotation error detection across eighteen methods and nine datasets, distinguishing between annotation errors, which are objectively incorrect, and disagreements that reflect plausible variation. No current method reliably separates the two phenomena without human adjudication. Building on this, Weber-Genzel et al. (2024) introduce the VariErr NLI dataset and a two-round annotation methodology for separating annotation errors from label variation, demonstrating that even state-of-the-art automatic methods fall short of human performance.

A complementary strand of work examines how the design stance of annotation guidelines shapes data quality. Röttger et al. (2022) distinguish *prescriptive* paradigms, which discourage annotator subjectivity to enable consistent training, from *descriptive* paradigms, which embrace it to model human diversity. Hate speech annotation sits un-

easily between these poles: prescriptive guidelines are necessary for consistency, yet culturally inadequate ones produce systematic error rather than consensus. Uma et al. (2021) survey methods for learning from disagreement, showing that structured disagreement patterns carry diagnostic value rather than being mere noise. Our work applies this insight to guideline design: the *structure* of disagreement localizes the failure mechanism.

Our contribution focuses on the upstream conditions that make certain errors structurally inevitable, prior to any annotator involvement. We position this against HLV in Section 5: the cases we describe are not legitimate variation to be modeled but underspecification or category mismatch to be repaired.

## 2.2 Hate Speech Annotation

Hate speech annotation is challenging because hate speech is context-dependent, culturally embedded, and subject to definitional disagreement across research communities, legal frameworks, and platform moderation policies (Waseem and Hovy, 2016; Davidson et al., 2017; Vidgen et al., 2019). Waseem and Hovy (2016) produced one of the earliest large-scale hate speech datasets, annotating English tweets within a framework grounded in Western civil rights categories. Davidson et al. (2017) extended this with a three-way distinction between hate speech, offensive language, and neither. Founta et al. (2018) showed that annotator pool composition substantially affects label distributions. Zampieri et al. (2020) organized a multilingual shared task and showed annotation schemes do not transfer reliably across languages. Vidgen et al. (2021) propose adversarial data collection as a remedy for dataset brittleness.

For Arabic specifically, Mubarak et al. (2017) developed an abusive language dataset using ternary classification, Albadi et al. (2018) constructed the first Arabic dataset for religious hate speech, and Mulki et al. (2019) introduced L-HSAB, a Levantine Twitter dataset. All three were developed with limited attention to dialectal variation and culturally specific insult categories, gaps this paper analyzes as sources of systematic guideline-induced error.

## 2.3 Arabic NLP and Annotation

The challenges of Arabic NLP are well documented. Habash (2010) surveys Arabic morphological complexity, orthographic ambiguity, and the

diglossic relationship between MSA and regional dialects. Darwish et al. (2017) demonstrate that Arabic POS tagging requires feature engineering tailored to Arabic-specific morphological phenomena, underscoring the difficulty of directly applying tools and frameworks developed for English. These challenges compound on social media, where non-standard spelling, code-switching, and high proportions of dialectal content create a particularly heterogeneous linguistic environment (Mubarak et al., 2017).

## 3 Arabic Social Media Discourse: Annotation Challenges

Arabic social media discourse presents four properties that make it particularly vulnerable to guideline-induced annotation errors. None of these properties is unique to Arabic in the strict sense: sarcasm and indirect insult exist in every language community, and many languages exhibit dialectal variation. The point below is more specific. The *combination* of these properties, together with specific categorical mismatches against dominant English language guidelines, produces a configuration of annotation risks for Arabic that has not been adequately documented.

### Dialectal variation as an annotation problem.

Arabic is a collection of related varieties in a diglossic relationship with MSA (Habash, 2010). Gulf, Levantine, Egyptian, and Maghrebi Arabic differ substantially at the lexical, morphological, and pragmatic levels (Darwish et al., 2017). On social media, users frequently mix dialects within a single post and use dialect as a marker of regional identity. This is a general fact about Arabic NLP, but it has a specific consequence for hate speech annotation: when a token is hostile in one dialect and neutral or even affectionate in another, an annotator’s dialect background biases the label they assign even when they apply the guideline faithfully. The result is disagreement that tracks annotators rather than texts, the diagnostic signature of guideline underspecification rather than legitimate label variation.

**Culturally embedded hostility.** Arabic hostile discourse frequently targets dimensions absent from English-language hate speech frameworks: lineage and ancestral origin (*nasab*, نَسَب *nasab*), tribal or clan affiliation, regional identity, and sectarian belonging (Mubarak et al., 2017; Albadi et al., 2018). The disparity between the categories

of social harm assumed by guidelines and those operative in the target discourse is the root cause of the errors analyzed here. Lineage and tribe based hostility do exist in Western contexts, but they rarely figure in the core categorical schemata of widely used guidelines, and so they become systematically invisible when those schemata are applied to Arabic data.

**Sarcasm and pragmatic indirectness.** Arabic online discourse makes extensive use of sarcasm, irony, and indirect criticism (Mulki et al., 2019). Hostile force depends on pragmatic inference rather than explicit surface content, so guidelines focused on surface features systematically misclassify sarcastic praise as neutral. While sarcasm is universal, what is specific here is the dense interaction between conventionalized honorific forms (such as *fandī*) and the pragmatic inversion that contexts of online conflict impose on them. Without explicit guidance on how to read these inversions, annotators have no shared procedure for resolving them.

**Religious and sectarian rhetoric.** Intra-Islamic sectarian hostility, denying the religious legitimacy of a specific sect, is a prevalent form of group-targeted harm in Arabic online discourse. Generic “religion” categories derived from English moderation policies do not distinguish this from anti-religious speech, creating a categorical gap that produces systematic misclassification (Albadi et al., 2018). As we discuss in Section 5, the boundary between sectarian hostility and ordinary theological disagreement is operationalizable through a small set of textual cues even though it is not trivial.

**Honor-based hostility: a closer look.** Because honor-linked insult patterns are not transparent to annotators outside the relevant cultural context, we expand briefly on the underlying social logic. In many Arabic-speaking communities, the social standing of a male addressee is publicly tied to the perceived conduct of his female relatives, a construct often referred to as *sharaf* (شَرَف *šarāf*). Hostile speech can therefore attack the addressee indirectly, by asserting or insinuating that a sister, mother, or wife behaves in ways the local norm system codes as dishonorable. The harm is double layered: it stigmatizes the woman through her imputed conduct and damages the addressee through the implication that he cannot uphold his familial honor. From an annotation perspective, surface-level hos-

tility detection misses these expressions because they typically contain no slur and no protected-attribute reference. The hostile force is recoverable only by an annotator who can read the implicit norm violation, and guidelines that do not surface this pattern make consistent annotation impossible regardless of annotator experience.

## 4 Three Mechanisms of Guideline-Induced Error

We now present each mechanism with illustrative case studies. The examples are drawn from attested patterns in Arabic social media discourse and are representative of instance types found in existing Arabic hate speech datasets (Mubarak et al., 2017; Albadi et al., 2018; Mulki et al., 2019).

### 4.1 Mechanism 1: Cultural Misclassification

Cultural misclassification occurs when guidelines define hostility exclusively in terms of categories from one cultural context, leaving annotators without adequate categories for culturally specific hostile expressions. Most hate speech schemes organize hostility around race, ethnicity, gender, sexual orientation, religion, and nationality (Davidson et al., 2017; Waseem and Hovy, 2016), reflecting Western civil rights frameworks that do not exhaust the dimensions of hostility in Arabic online discourse.

**Case Study 1: Lineage insults.** In many Arabic-speaking communities, lineage (*nasab*, نَسَب *nasab*) carries deep social meaning. Attacking a person’s family origin by implying they lack honorable ancestry is a well-established form of hostile expression. The following attested type is common in Gulf Arabic social media:

**Arabic:** يَا قَدِيي الْأَصْل *yā qadiyy al-ʾaṣl*

**Gloss:** “O you of ignoble/lowly origin”

**Note:** Attacks the addressee’s family lineage and social standing; highly offensive in Gulf and Levantine contexts.

Table 1 shows a representative annotation pattern for this instance type across three annotators working with a standard hate-speech guideline.

The inconsistency in Table 1 is not annotator error. It arises because the guidelines provide no category for ancestry-based hostility. Annotator B treats lineage as a proxy for social group

Annotator	Label	Rationale
A	insult	personal attack on addressee
B	hate speech	targets a social group
C	offensive, not hate	no protected category matched

Table 1: Annotation disagreement for a lineage insult. The inconsistency arises because guidelines provide no category for ancestry-based hostility, not because annotators have erred.

membership; Annotator C, correctly noting no protected category matches as written, assigns a lower-severity label. Both act in conformity with the guidelines. The error is structural.

**Case Study 2: Tribal and regional insults.** Insults targeting tribal affiliation or regional origin carry significant social force in Gulf and Levantine contexts. Consider the following example type:

**Arabic:** أَنْتَ مِنْ قَبِيلَتِهِمُ الْمُتَخَلِّفَهُ *anta min qabiylatihim al-mutahālifah*

**Gloss:** “You are from their backward tribe”

**Note:** Stigmatizes a named or implied tribal group; functions as group-targeted hate speech in social context but falls outside standard annotation categories.

Such expressions imply membership in a stigmatized tribe or region, functioning as group-targeted hate speech. Annotation schemes that omit regional or within-community ethnicity as protected dimensions produce inconsistent labels that annotator training alone cannot resolve without prior guideline revision.

#### 4.2 Mechanism 2: Dialectal Ambiguity

Dialectal ambiguity arises when guidelines assume stable lexical semantics across Arabic varieties. Many lexical items vary substantially in pragmatic force across dialects, so guidelines without dialect-sensitive examples produce annotation errors even when annotators follow instructions correctly.

**Case Study 3: The word *ḥayawān*.** The Arabic word for “animal” (*ḥayawān*, حَيَوَان *ḥayawān*) illustrates dialectal ambiguity. In MSA and formal contexts it is a neutral descriptor. Its pragmatic force varies considerably across dialectal and social contexts, as shown in Table 2.

Context	Variety	Pragmatic Force
descriptive statement	MSA	neutral
hostile comment	Gulf colloquial	strong insult
teasing among friends	Egyptian colloquial	affectionate banter
online argument	Levantine colloquial	mild to strong insult

Table 2: Dialectal and contextual variation in the pragmatic force of *ḥayawān* (حَيَوَان *ḥayawān*) across Arabic varieties.

Example A is hostile; Example B is not. A guideline that treats *ḥayawān* as inherently offensive will mislabel Example B. One that treats it as context-dependent but provides no dialect-sensitive guidance produces label variation correlated with annotator dialect rather than post content. This is guideline underspecification, not genuine HLV: the two examples are fully resolvable for annotators who share the relevant dialect. The following attested pair illustrates the divergence concretely:

**Example A** (hostile, Gulf Arabic):

أَنْتَ حَيَوَانٌ وَلَا تَسْتَجِلُّ الرَّدَّ *anta ḥayawān walā tistaḥil al-radd*

“You are an animal and do not deserve a response.”

**Example B** (affectionate banter, Egyptian Arabic):

إِنَّتَ حَيَوَانٌ يَا ابْنَ، كَيْدَ بِيَهْبُونِي *enta ḥayawān ya 'bni, keda bithibbuniy*

“You animal, my son, is that how you show you care?”

**Case Study 4: Honorifics used sarcastically.** In Levantine and Gulf Arabic, honorific address forms that are conventionally respectful can be deployed sarcastically to convey contempt. The highly subjective nature of Arabic sarcasm annotation, and its dependence on dialect and cultural context, has been empirically documented: [Abu Farha and Magdy \(2020\)](#) show that annotator dialect background substantially shifts sarcasm labels in Arabic social media data. Consider the following example:

**Arabic:** تَبَارَكَ اللَّهُ عَلَيْكَ يَا فَنْدِي، مَعْرِفَتَكَ بَتُحْرِفُ *tabaraka allah alayka ya fandiy, marifatak btuhrif*

**Literal gloss:** “God bless you, sir, your knowledge is enlightening.”

**Pragmatic force:** Contemptuous sarcasm; the honorific *fandī* and the apparent praise are used ironically to ridicule the addressee’s ignorance.

Unlike the *ḥayawān* case, the error here is not a matter of which meaning is primary but of whether pragmatic inversion is recognized as a category of hostile expression at all. Guidelines that do not address sarcasm and do not provide dialect-specific examples of this pattern produce systematic misclassifications. We acknowledge that sarcasm cases sit closest to genuine HLV among the patterns we describe: in some posts, sarcastic and sincere readings are both available, and disagreement reflects real interpretive plurality. Our claim is narrower: even setting those genuinely ambiguous cases aside, a substantial residue of disagreement remains that is attributable to the guideline never naming pragmatic inversion as a category to track.

### 4.3 Mechanism 3: Annotation Projection

Annotation projection refers to transferring annotation schemes developed for one language or cultural context to a substantially different one without adequate adaptation. In Arabic hate speech annotation this most commonly manifests as English-language moderation categories applied to Arabic discourse, introducing categorical mismatches annotators cannot resolve within the scheme.

**Case Study 5: Sectarian hostility.** Consider the following attested expression type from Arabic social media:

**Arabic:** هَاؤُلَاءِ لَا يَنْتَمُونَ إِلَى الْإِسْلَامِ الْحَقِيقِيِّ  
*hā’ulā’ lā yantamūna ilā ‘l-islām al-ḥaqiqiy*

**Gloss:** “These people do not belong to true Islam.”

**Note:** Targets members of a specific Islamic sect; a form of religiously motivated group targeting with severe social consequences in many Arabic-speaking contexts.

Many frameworks derived from English moderation policies collapse sectarian speech into a generic “religion” category that does not distinguish anti-religion from intra-religion hostility, or lack any category for intra-religious targeting. The result aligns with the cross-lingual transfer challenge documented by [Zampieri et al. \(2020\)](#): an-

notation schemes do not transfer reliably across languages.

**Case Study 6: Gender and honor-linked hostility.** English-language hate speech frameworks typically treat gender-based hostility in terms of misogyny or sexism directed against women as a social group. Arabic online discourse, however, includes forms of gendered insult structured around the concept of *sharaf* (honor, شَرَف *šarāf*). Consider the following example:

**Arabic:** أُخْتُكَ بِتَسْمَرٍ بَرَّةٍ بِلَيْلٍ *uhtuka bitsa-har barrah bil-layl*

**Gloss:** “Your sister stays out late at night.”

**Note:** An honor-based attack on the male addressee through impugning a female relative’s behavior; culturally potent as a hostile move in Gulf and Levantine contexts.

This form of hostility is simultaneously gendered and honor-linked in ways that do not map onto the sexism or misogyny categories of English-derived frameworks: harm targets both the woman discussed and the male addressee whose honor is attacked through association. Guidelines that import English-derived gendered hostility categories without adaptation will systematically under-detect this culturally prevalent form of harm.

## 5 Taxonomy and Diagnostic Framework

Based on the analysis in Section 4, we propose a taxonomy of guideline-induced annotation errors for Arabic hate speech corpora. The five error types in Table 3 are distinguished by their primary cause in guideline design and by the characteristic misannotation pattern each produces.

This taxonomy is explicitly distinct from human label variation in the sense of [Weber-Genzel et al. \(2024\)](#) and the broader HLV literature ([Plank et al., 2014](#); [Pavlick and Kwiatkowski, 2019](#); [Nie et al., 2020](#); [Jiang et al., 2023](#)). The errors we describe are not cases of genuine ambiguity: the guideline itself provides inadequate or misleading instruction, making consistent application impossible regardless of annotator competence. Where HLV may be addressed by modeling disagreement ([Uma et al., 2021](#)), guideline-induced errors require upstream revision of the annotation framework. Crucially, the *structure* of observed disagreement signals which revision is needed: clustering by annotator dialect indicates dialectal ambiguity; clustering

Error Type	Guideline Cause	Manifestation	Arabic Example
Cultural misclassification	Category set excludes culturally relevant hostility dimensions	Inconsistent labeling of expressions targeting lineage, tribe, or regional identity	يَا قَدِي الْأَصْلِ <i>yā-qadiyy al-aṣl</i>
Dialectal ambiguity	Guidelines assume uniform lexical semantics across varieties	Same token receives different labels from annotators of different dialect backgrounds	حَيَوَانِ <i>ḥayawān</i> as insult vs. banter
Annotation projection	English-derived categories transferred without adaptation	Sectarian speech collapsed into generic religion label; honor-linked gendered insults misclassified	هَآؤُلَاءِ لَا يَنْتَمُونَ <i>hā-ulā lā yantamūna</i> ; أُخْتُكَ بِتَسَهَّرَ <i>uhtuka bitsahar</i>
Pragmatic misinterpretation	Guidelines do not address sarcasm or indirect hostility	Sarcastic praise annotated as neutral; ironic honorifics labeled as non-offensive	تَبَرَّكَ آلَهُ عَلَيْهِ بِي فَندِي <i>tabaraka āllah ‘alayka ya fandiy</i>
Boundary underspecification	Annotation span boundaries undefined for multi-unit expressions	Inconsistent annotation of multi-token or idiomatic insult constructions	Multi-word dialect insult idioms

Table 3: Taxonomy of guideline-induced annotation errors in Arabic hate speech corpora. Error types are distinguished by their upstream cause in guideline design rather than by annotator behavior.

by cultural background indicates cultural misclassification; uniform disagreement suggests boundary underspecification.

### 5.1 Mapping Mechanisms to Error Types

The relationship between the three mechanisms in Section 4 and the five error types in Table 3 is many-to-many rather than one-to-one, made explicit in Table 4. Cultural misclassification as a mechanism produces the cultural misclassification error type directly and contributes to annotation projection errors when an imported scheme imposes the wrong categorical fit. Dialectal ambiguity produces both the dialectal ambiguity error type and a share of pragmatic misinterpretation errors when dialect licenses the pragmatic inversion. Annotation projection feeds into several error types simultaneously. Boundary underspecification cuts across all three mechanisms whenever the guideline does not specify how to delimit multi-unit hostile expressions.

### 5.2 Extended Category Schema for Arabic

To operationalize the cultural misclassification and annotation projection error types, we propose five culturally grounded hostility dimensions that standard English-derived guidelines omit. For each, we provide a working definition and a positive/negative example pair.

**Lineage hostility** targets a person’s ancestral origin (*nasab*, نَسَب *nasab*). *Positive*: explicit

Mechanism	Error types it primarily produces
Cultural misclassification	Cultural misclassification (primary); annotation projection (secondary)
Dialectal ambiguity	Dialectal ambiguity (primary); pragmatic misinterpretation (secondary, via sarcasm)
Annotation projection	Annotation projection (primary); cultural misclassification (secondary); pragmatic misinterpretation (secondary)
All three	Boundary underspecification (whenever multi-token expressions are involved)

Table 4: Mapping from mechanisms (Section 4) to error types (Table 3). The relationship is many-to-many: a single mechanism can drive several error types, and a single error type can arise from more than one mechanism.

imputation of ignoble ancestry (*yā qadīy al-aṣl*). *Negative*: statements about genealogy that do not stigmatize.

**Tribal/regional hostility** stigmatizes membership in a named or implied tribal, clan, or regional group. *Positive*: attributing backwardness or inferiority to a named tribe or region. *Negative*: neutral identification of regional origin without derogatory framing.

**Sectarian hostility** targets members of a specific religious sub-community through exclusion or delegitimization. *Positive*: denying the Islamic cre-

dentials of a sectarian group (*hā’ulā’ lā yantamūna ilā al-islām al-ḥaḳīqī*). *Negative*: neutral theological disagreement without group-targeting framing. The boundary between sectarian hostility and theological disagreement, while non-trivial, admits an operational test based on three textual cues: (i) the predicate targets the *persons* affiliated with the sect rather than a doctrinal proposition; (ii) the framing involves exclusion from a religiously legitimate in-group (formulations such as “not real Muslims”, “outside true Islam”); and (iii) the expression carries an exhortative or stigmatizing illocutionary force rather than a propositional one. Disagreements with a doctrine in the abstract, without these features, fall outside the category.

**Honor-linked gendered hostility** attacks a male addressee through the implied behavior of a female relative. *Positive*: insinuating that a sister or mother violates honor norms (*ukhtuka bitsahar barra bil-layl*). *Negative*: neutral reference to a female relative’s activities without honor-impugning implicature.

**Pragmatic inversion** applies to sarcasm and ironic honorifics where surface praise encodes contempt. *Positive*: honorific address combined with contextual cues of ridicule (*tabaraka Allāh ‘alayka ya fandī*). *Negative*: sincere use of the same honorific. Where expressions encode multiple dimensions simultaneously, annotators should apply multi-labeling. A priority rule is needed only for span-level annotation of the *primary target of harm*: the directly addressed party takes precedence over secondary targets.

### 5.3 Diagnostic Framework

We propose a four-stage process for identifying guideline-induced errors during corpus construction, designed as an actionable checklist for dataset builders. The stages apply iteratively, with later findings feeding back into earlier revisions.

## 6 Implications

### 6.1 For Annotation Practice

Annotation guidelines should be developed by teams combining computational linguistics expertise with cultural knowledge of Arabic-speaking communities (Vidgen et al., 2019). The choice between prescriptive and descriptive paradigms (Röttger et al., 2022) should be made explicitly: prescriptive guidelines are appropriate for consistent moderation but must be prescriptive

Stage	Key diagnostic question
<b>1. Cultural audit</b>	Does the guideline cover hostility targeting lineage, tribe, region, or sect? Are community members and cultural linguists consulted?
<b>2. Dialect calibration</b>	Does within-dialect $\kappa$ substantially exceed overall $\kappa$ ? Are dialect-sensitive examples provided for known ambiguous items (e.g., <i>ḥayawān</i> , <i>majnūn</i> )? Use CAMELira (Obeid et al., 2022) for disambiguation; Kumar (Khalifa et al., 2016) for Gulf examples.
<b>3. Projection check</b>	Were categories imported from another language? Do they cover sectarian targeting and honor-linked gendered hostility (Mubarak et al., 2017; Albadi et al., 2018)? Categories with low pilot coverage must be revised.
<b>4. Pragmatic specification</b>	Does the guideline address sarcasm and ironic honorifics (Röttger et al., 2022)? Is a SARCASM-FLAG label with rationale fields required? Persistent post-training disagreement signals a category gap (Swayamdipta et al., 2020; Abu Farha and Magdy, 2020; Darwish et al., 2017).

Table 5: Diagnostic framework for guideline-induced errors. Each stage targets one error type from the taxonomy and provides a concrete decision criterion for guideline revision.

about *Arabic-specific* hostility dimensions. Annotators should be recruited with known dialect backgrounds, since within-dialect agreement substantially exceeding cross-dialect agreement signals a candidate item set for guideline revision rather than adjudication. Following Swayamdipta et al. (2020), pilot studies should diagnose guideline failures, not only estimate IAA: the disagreement correlation structure localizes the mechanism.

### 6.2 For NLP Systems

Guideline-induced errors propagate directly into trained models. Models built on datasets with systematic cultural misclassifications will under-detect culturally specific hostility; those trained on dialectally ambiguous labels may learn spurious form-to-label correlations that fail across dialect boundaries. Adversarial data collection (Vidgen et al., 2021) addresses symptoms rather than causes: categorical gaps cannot be closed by collecting more examples within the same scheme. The errors we describe are not random noise that larger datasets will average out; they are systematic biases that larger datasets will amplify.

### 6.3 For Dataset Documentation

Dataset papers should include a *guideline scope statement* specifying which cultural dimensions of hostility the annotation scheme is and is not designed to capture. Users of corpora such as Mubarak et al. (2017), Albadi et al. (2018), and Mulki et al. (2019) need to know that models trained on these datasets may under-detect lineage-based, sectarian, or honor-linked hostility, not because the datasets are flawed in execution but because their guidelines were not scoped to cover these dimensions. The cultural adequacy audit (Stage 1) requires community participation beyond standard crowdsourcing (Vidgen et al., 2019).

## 7 Related Work

Klie et al. (2023) argue that annotation quality requires both detecting errors in existing datasets and preventing them in new ones; our paper addresses prevention via a diagnostic framework for guideline-induced errors. The prescriptive versus descriptive paradigm distinction (Röttger et al., 2022) is directly relevant: prescriptive guidelines are necessary but produce systematic error when culturally misaligned with the target discourse. Uma et al. (2021) show that disagreement patterns carry diagnostic value; our framework operationalizes this to localize the active error mechanism. The broader HLV literature (Plank et al., 2014; Pavlick and Kwiatkowski, 2019; Nie et al., 2020; Jiang et al., 2023; Giulianelli et al., 2023; Weber-Genzel et al., 2024) provides the conceptual scaffolding within which our work is positioned: it argues for taking disagreement seriously as a signal rather than noise, and we extend that argument by carving out a class of disagreement that signals upstream guideline failure rather than downstream interpretive plurality.

Research on Arabic hate speech has grown substantially. Mubarak et al. (2017), Albadi et al. (2018), and Mulki et al. (2019) established foundational datasets whose annotation guidelines have not been systematically examined for the mechanisms we describe. Haj Ahmed et al. (2025) show that L-HSAB exhibits a strong Lebanese-dialect bias, causing models to generalize poorly across other Levantine varieties: a concrete instance of dialectal ambiguity and annotation projection combining to produce dataset brittleness. The need for culturally grounded categories extends beyond hate speech. Al-Khalifa et al. (2026) introduce ADAB,

a large-scale Arabic politeness dataset covering MSA and four dialect groups annotated according to Arabic linguistic traditions and pragmatic theory, demonstrating the viability of Arabic-specific annotation schemata grounded in sociopragmatic norms. Abu Farha and Magdy (2020) demonstrate that annotator dialect background shifts sarcasm labels in Arabic social media, providing direct evidence for the dialectal ambiguity mechanism and supporting our Stage 2 recommendation. Zampieri et al. (2020) document cross-lingual transfer failures, and Founta et al. (2018) show that annotator pool composition substantially affects label distributions; our contribution is demonstrating that disagreement *structure* also localizes the active taxonomy error type.

While our analysis focuses on Arabic, the mechanisms are not in principle Arabic-specific. Any language whose discourse community exhibits a diglossic structure, a dialect continuum, or culturally embedded hostility lacking equivalents in dominant English-derived frameworks is in principle exposed to analogous risks. We do not have the empirical basis here to demonstrate parallel mechanisms in specific other languages, and we therefore present this generalization as a conjecture for future work rather than a demonstrated finding.

## 8 Conclusion

This paper argues that a substantial class of annotation errors in Arabic hate speech corpora originates from structural weaknesses in guideline design rather than annotator failure. We identify three mechanisms, cultural misclassification, dialectal ambiguity, and annotation projection, and show through six case studies how each produces systematic errors distinct from normal label variation. We propose a five-type taxonomy, an explicit mapping from mechanisms to error types, and a four-stage diagnostic framework instantiated for Arabic with specific lexical items and expression types. The paper reframes annotation quality as a problem of guideline design rather than annotator behavior, intervening before annotation begins by targeting upstream sources of error. For Arabic and other languages whose discursive cultures are poorly represented in existing frameworks, this reframing is crucial for building NLP systems that are both accurate and fair.

## Limitations

The case studies presented in this paper are illustrative rather than empirical: they demonstrate the plausibility and linguistic basis of the error mechanisms we identify but do not constitute a quantitative analysis of error rates in specific datasets. We therefore frame the present contribution as theoretical and programmatic rather than empirical. The diagnostic framework in Section 5 has not yet been applied end-to-end to a complete dataset, and the operational thresholds we suggest (such as comparing within-dialect to overall agreement) need calibration through use. Reviewers of the present version raised this point pointedly, and we agree it is the most important next step.

A full empirical evaluation would require access to complete annotation logs with annotator-level data and dialect metadata for existing Arabic hate speech datasets, which is not uniformly available in the current literature. Future work should combine the diagnostic framework proposed here with the annotation error detection methods surveyed in Klie et al. (2023) to produce quantitative estimates of error rates attributable to each mechanism. A controlled pilot study applying Stages 1–4 to a subset of an existing dataset such as L-HSAB (Mulki et al., 2019) would allow before-and-after measurement of IAA and dialect-conditioned agreement, providing concrete evidence for the gains our framework is designed to produce. We regard the present work as providing the theoretical grounding and operational vocabulary that such a study requires.

The taxonomy we propose is based on analysis of Arabic social media discourse and may not be exhaustive. Additional mechanisms of guideline-induced error may exist for Arabic or for other languages with similarly complex sociolinguistic profiles. The generalization to other diglossic or low-resource languages, made cautiously in Section 7, has not been validated empirically and should be read as a conjecture inviting further work rather than as a demonstrated finding.

Finally, although we contrast guideline-induced error with HLV throughout, we acknowledge that in practice the two coexist: a single disagreement event may reflect both genuine interpretive plurality and underspecified guideline categories. Disentangling them at the instance level remains an open problem, and our framework is meant to complement rather than replace methods that target HLV

directly.

## Ethical Considerations

This paper discusses hate speech and abusive language in Arabic. All examples are drawn from or modeled on attested discourse phenomena in Arabic social media, and none are fabricated for the purpose of demonstrating harmful content. We do not release any datasets or annotation resources as part of this submission.

The diagnostic framework proposed here is intended to improve the quality of hate speech annotation guidelines, enabling more accurate and fair automated detection of harmful content. We encourage researchers who adopt our framework to consult community stakeholders during the cultural adequacy audit stage and to consider the downstream uses of the corpora they construct.

The Arabic linguistic examples in this paper are used analytically and are not intended to propagate or endorse the expressions they exemplify. All examples are presented alongside translations and pragmatic notes to prevent out-of-context misuse.

## Acknowledgment

This work was made possible by the National Priorities Research Program grant NPRP14C-0916-210015 from the Qatar Development and Innovation Council (QRDI).

## References

- Ibrahim Abu Farha and Walid Magdy. 2020. From Arabic sentiment analysis to sarcasm detection: The ArSarcasm dataset. In *Proceedings of the 4th Workshop on Open-Source Arabic Corpora and Processing Tools, with a Shared Task on Offensive Language Detection*, pages 32–39, Marseille, France. European Language Resource Association.
- Nuha Albadi, Maram Kurdi, and Shivakant Mishra. 2018. Are they our brothers? Analysis and detection of religious hate speech in the Arabic Twittersphere. In *Proceedings of the 2018 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*, pages 69–76. ACM.
- Hend Al-Khalifa, Nadia Ghezaiel, Maria Bounnit, Hend Hamed Alhazmi, Noof Abdullah Alfear, Reem Fahad Alqifari, Ameer Masoud Almasoud, and Sharefah Ahmed Al-Ghamdi. 2026. ADAB: Arabic dataset for automated politeness benchmarking, a large-scale resource for computational sociopragmatics. In *Proceedings of the Fifteenth biennial Language Resources and Evaluation Conference (LREC)*

- 2026), Palma, Mallorca, Spain. European Language Resources Association. To appear.
- Ron Artstein and Massimo Poesio. 2008. Survey article: Inter-coder agreement for computational linguistics. *Computational Linguistics*, 34(4):555–596.
- Kareem Darwish, Hamdy Mubarak, Ahmed Abdelali, and Mohamed Eldesouki. 2017. Arabic POS tagging: Don’t abandon feature engineering just yet. In *Proceedings of the Third Arabic Natural Language Processing Workshop*, pages 130–137. Association for Computational Linguistics.
- Thomas Davidson, Dana Warmusley, Michael Macy, and Ingmar Weber. 2017. Automated hate speech detection and the problem of offensive language. In *Proceedings of the 11th International AAAI Conference on Web and Social Media (ICWSM)*, pages 512–515.
- Antigoni-Maria Founta, Constantinos Djouvas, Despoina Chatzakou, Ilias Leontiadis, Jeremy Blackburn, Gianluca Stringhini, Athena Vakali, Michael Sirivianos, and Nicolas Kourtellis. 2018. Large scale crowdsourcing and characterization of Twitter abusive behavior. In *Proceedings of the 12th International AAAI Conference on Web and Social Media (ICWSM)*.
- Mario Giulianelli, Joris Baan, Wilker Aziz, Raquel Fernández, and Barbara Plank. 2023. What comes next? Evaluating uncertainty in neural text generators against human production variability. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 14349–14371, Singapore. Association for Computational Linguistics.
- Nizar Habash. 2010. *Introduction to Arabic Natural Language Processing*. Morgan & Claypool Publishers.
- Ahmed Haj Ahmed, Rui-Jie Yew, Xerxes Minocher, and Suresh Venkatasubramanian. 2025. Navigating dialectal bias and ethical complexities in Levantine Arabic hate speech detection. In *Proceedings of the 4th Workshop on Arabic Corpus Linguistics (WACL-4)*, pages 103–108, Abu Dhabi, UAE. Association for Computational Linguistics.
- Nan-Jiang Jiang, Chenhao Tan, and Marie-Catherine de Marneffe. 2023. Ecologically valid explanations for label variation in NLI. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 10622–10633, Singapore. Association for Computational Linguistics.
- Salam Khalifa, Nizar Habash, Dana Abdulrahim, and Sara Hassan. 2016. A large scale corpus of Gulf Arabic. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC 2016)*, pages 4282–4289.
- Jan-Christoph Klie, Bonnie Webber, and Iryna Gurevych. 2023. Annotation error detection: Analyzing the past and present for a more coherent future. *Computational Linguistics*, 49(1):157–198.
- Hamdy Mubarak, Kareem Darwish, and Walid Magdy. 2017. Abusive language detection on Arabic social media. In *Proceedings of the First Workshop on Abusive Language Online*, pages 52–56. Association for Computational Linguistics.
- Hala Mulki, Hatem Haddad, Chedi Bechikh Ali, and Halima Alshabani. 2019. L-HSAB: A Levantine Twitter dataset for hate speech and abusive language. In *Proceedings of the Third Workshop on Abusive Language Online*, pages 111–118. Association for Computational Linguistics.
- Yixin Nie, Xiang Zhou, and Mohit Bansal. 2020. What can we learn from collective human opinions on natural language inference data? In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 9131–9143, Online. Association for Computational Linguistics.
- Ossama Obeid, Go Inoue, and Nizar Habash. 2022. CAMELira: An Arabic multi-dialect morphological disambiguator. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 319–326, Abu Dhabi, UAE. Association for Computational Linguistics.
- Ellie Pavlick and Tom Kwiatkowski. 2019. Inherent disagreements in human textual inferences. *Transactions of the Association for Computational Linguistics*, 7:677–694.
- Barbara Plank, Dirk Hovy, and Anders Søgaard. 2014. Linguistically debatable or just plain wrong? In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 507–511. Association for Computational Linguistics.
- Paul Röttger, Bertie Vidgen, Dirk Hovy, and Janet Pierrehumbert. 2022. Two contrasting data annotation paradigms for subjective NLP tasks. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 175–190, Seattle, United States. Association for Computational Linguistics.
- Swabha Swayamdipta, Roy Schwartz, Nicholas Lourie, Yizhong Wang, Hannaneh Hajishirzi, Noah A. Smith, and Yejin Choi. 2020. Dataset cartography: Mapping and diagnosing datasets with training dynamics. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 9275–9293. Association for Computational Linguistics.
- Alexandra N. Uma, Tommaso Fornaciari, Dirk Hovy, Silviu Paun, Barbara Plank, and Massimo Poesio. 2021. Learning from disagreement: A survey. *Journal of Artificial Intelligence Research*, 72:1385–1470.
- Bertie Vidgen, Alex Harris, Dong Nguyen, Rebekah Tromble, Scott Hale, and Helen Margetts. 2019.

- Challenges and frontiers in abusive content detection. In *Proceedings of the Third Workshop on Abusive Language Online*, pages 80–99. Association for Computational Linguistics.
- Bertie Vidgen, Trista Thrush, Zeerak Waseem, and Douwe Kiela. 2021. Learning from the worst: Dynamically generated datasets to improve online hate detection. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1667–1682. Association for Computational Linguistics.
- Zeerak Waseem and Dirk Hovy. 2016. Hateful symbols or hateful people? Predictive features for hate speech detection on Twitter. In *Proceedings of the NAACL Student Research Workshop*, pages 88–93. Association for Computational Linguistics.
- Leon Weber-Genzel, Siyao Peng, Marie-Catherine de Marneffe, and Barbara Plank. 2024. VariErr NLI: Separating annotation error from human label variation. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2256–2269. Association for Computational Linguistics.
- Marcos Zampieri, Preslav Nakov, Sara Rosenthal, Pepa Atanasova, Georgi Karadzhov, Hamdy Mubarak, Leon Derczynski, Zeses Pitenis, and Çağrı Çöltekin. 2020. SemEval-2020 task 12: Multilingual offensive language identification in social media (Off-ComEval). In *Proceedings of the Fourteenth Workshop on Semantic Evaluation*, pages 1425–1447. International Committee for Computational Linguistics.

# When LLMs Disagree with Human Experts: Understanding LLM Annotation Failures in Nutrition Misinformation through Hierarchical Error Analysis using Seed Oil Narratives

**Vishwaa Shah**

University of North Florida  
School of Computing  
Jacksonville, FL, USA  
N01458714@unf.edu

**Indika Kahanda**

University of North Florida  
School of Computing  
Jacksonville, FL, USA  
indika.kahanda@unf.edu

**Andrea Arikawa**

University of North Florida  
Nutrition & Dietetics  
Jacksonville, FL, USA  
a.arikawa@unf.edu

## Abstract

Accurate linguistic annotation is crucial for creating high-quality datasets in specialized domains, yet manual labeling is often slow, expensive, and inconsistent. We present a reproducible workflow for evaluating the effectiveness of large language models (LLMs) as annotators of domain-specific health misinformation on social media. Using a data set of 169 Instagram posts on seed oils, expert nutritionists provided gold-standard labels (71% positives), which we compared against the outputs of five open-source LLMs. We introduce a hierarchical error taxonomy that categorizes LLM misclassifications according to the direction, mechanism, and contributing factors of the error, providing interpretable insights into model failures. Our analysis reveals systematic error patterns, including misinterpretation of nuanced claims and overconfidence in predictions, highlighting conditions under which LLM annotations do not align with expert judgment. Although the data set is modest in size and exhibits class imbalance, it reflects real-world distributions of nutrition-related Instagram content and motivates the need for a careful evaluation of the robustness of the LLM annotation. This study has implications for the development of frameworks for automated LLM-based annotators in the health and nutrition domains, as well as LLM developers in general.

## 1 Introduction

Nutrition misinformation on social media poses a growing public health challenge, shaping perceptions and behaviors around diet and health (Faruk, 2024). Platforms like Instagram, with over two billion monthly users (Dean, 2025), amplify the spread of such content due to their visual and influence-driven ecosystems. A prominent example is discourse around “seed oils,” commonly consumed vegetable oils rich in unsaturated fatty acids. Although the U.S. Dietary Guidelines recommend

their inclusion in a healthy diet (Dietary Guidelines for Americans, 2020), online narratives frequently portray these oils as harmful or inflammatory, linking them to chronic diseases (Petersen et al., 2024). Surveys indicate that 43% of U.S. consumers recently encountered information about seed oils primarily via social media rather than healthcare professionals or scientific sources (Balagtas and Bryant, 2025).

Given the scale and speed of misinformation, automated detection methods are increasingly important. LLMs such as GPT-4, ChatGPT, and Llama have demonstrated strong capabilities in identifying misleading health-related content (Tan et al., 2025) and can also serve as annotators, generating labels more efficiently than human experts (Tan et al., 2024; Goel et al., 2023). However, domain-specific claims, such as those regarding seed oils or n-6 PUFA intake, pose additional challenges due to complex and nuanced scientific evidence (Petersen et al., 2024).

Recent approaches improve annotation quality by combining multiple LLMs or leveraging human-in-the-loop strategies, including ensemble methods, relevancy scoring, and multi-step verification frameworks (Qiu et al., 2025; Schroeder et al., 2025). Human-LLM collaborative frameworks further enhance reliability by using model explanations to guide selective re-annotation (Wang et al., 2024; Nahum et al., 2025). In the nutrition domain, fine-tuned models like BERT and RoBERTa have successfully detected misleading content on Instagram (Lamichhane et al., 2025), highlighting the need for robust annotation pipelines that align automated detection with expert judgment (Segado-Fernández et al., 2025).

Despite these advances, systematic evaluation of LLMs as domain-specific annotators remains limited. Misperceptions about seed oils persist online, and methods for integrating LLM outputs with high-quality expert annotation are underex-

plored. In this study, we evaluate multiple LLMs on Instagram captions about nutrition using the U.S. Dietary Guidelines as a normative reference. We compare model outputs to expert nutritionist labels and analyze annotation errors to identify common failure modes, providing actionable guidance for incorporating LLMs into annotation workflows.

While our study focuses on seed oils, it serves as a case study for evaluating LLM-based annotation workflows in a domain-specific setting. The proposed framework is designed to be adaptable to other annotation tasks, although it is empirically validated only in the nutrition misinformation domain. This framework captures linguistic and semantic nuances, claim hedging, and contextual interpretation, providing insights for designing reliable annotation processes across domains. We frame this study as an annotation error analysis, showing how evaluating human-machine disagreements can inform best practices for high-quality, domain-specific datasets. Because our operational definition of scientific consensus relies exclusively on the U.S. Dietary Guidelines Report, the resulting annotations are grounded in a U.S.-centric nutrition framework. This choice ensures consistency across expert coders but also narrows the scope of generalizability to other cultural or dietary standards.

Our contributions are threefold: (1) Present a reproducible workflow for evaluating LLMs as annotators of domain-specific health misinformation, (2) Introduce a hierarchical error taxonomy that categorizes LLM annotation failures by direction, mechanism, and contributing factors, and (3) Provide an empirical analysis of LLM annotation behavior on nutrition misinformation, identifying systematic error patterns that inform best practices for integrating LLMs into linguistic annotation pipelines.

## 2 Related Work

Previous work demonstrates that LLMs such as ChatGPT, GPT-4, and LLama can identify misleading health-related information with promising performance (Tan et al., 2025; Faruk, 2024; Yeung et al., 2022). These studies highlight the potential of LLMs for automated fact-checking, but their effectiveness varies across specialized topics, platforms, and cultural contexts, leaving certain domains, such as diet-specific misinformation, underexplored.

Surveys and empirical studies show that LLM-generated annotations can accelerate dataset cre-

ation, improve consistency, and reduce annotation costs (Tan et al., 2024; Goel et al., 2023; Alizadeh et al., 2025). Approaches like ensemble methods and relevancy scoring further enhance annotation reliability, mitigating heterogeneity in labeling decisions across multiple models (Qiu et al., 2025; Schroeder et al., 2025). Additionally, LLMs have been shown to detect label errors in existing datasets, improving model evaluation and downstream performance (Nahum et al., 2025).

Despite these advances, research has largely focused on generic misinformation or general annotation tasks. There remains a gap in applying LLM-based detection and annotation to domain-specific misinformation, such as dietary claims about seed oils. Our hierarchical error taxonomy extends prior work on LLM-assisted annotation by integrating direction, mechanism, and contextual contributing factors into a single framework. Unlike existing taxonomies that focus primarily on factuality or reasoning alone (Tan et al., 2024; Goel et al., 2023; Nahum et al., 2025), this structure additionally captures pragmatic and linguistic sources of misalignment such as claim hedging, evidence cues, and contextual interpretation providing a more comprehensive lens for analyzing annotation failures (Wang et al., 2024; Schroeder et al., 2025).

## 3 Methodology

Our methodology, shown in Figure 1, follows a four-stage pipeline. First, we collected publicly available Instagram posts discussing seed oils and preprocessed captions for textual analysis. Second, three nutrition experts independently annotated each post as *Credible* or *Misinformation* according to the 2020–2025 U.S. Dietary Guidelines, producing a gold-standard dataset. Third, five instruction-tuned LLMs were prompted using a knowledge-augmented Chain-of-Thought (CoT) approach to generate step-by-step reasoning, binary labels, and confidence scores for each post. Finally, LLM predictions were evaluated against expert labels using standard classification metrics, and misclassifications were analyzed using a hierarchical error taxonomy to identify systematic error patterns and contextual factors that affect model performance.

### 3.1 Data Collection

A nutrition expert identified English-language posts using Instagram’s native search functionality to locate content discussing dietary fats. Only

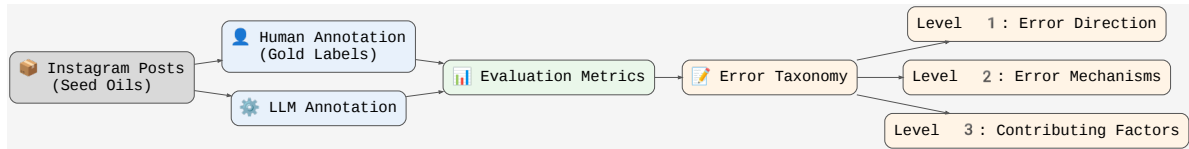


Figure 1: Full methodology pipeline from data collection to hierarchical error taxonomy analysis. The figure illustrates all major steps in our workflow, providing a visual guide to the complete processes.

posts with substantive discussion of seed oils were retained. Posts were considered relevant if captions explicitly referenced seed oils or specific seed oil products (e.g., soybean, sunflower, safflower, canola, or corn oil) in the context of nutritional or health claims. Seed oils were selected for analysis because they are frequently misrepresented on social media as harmful or inflammatory, despite being recommended sources of unsaturated fats by the U.S. Dietary Guidelines for Americans 2020–2025 (Dietary Guidelines for Americans, 2020).

To reduce personalization and geographic bias, searches were conducted using newly created, clean accounts with no prior activity, and all sessions used a VPN. This ensured retrieved posts reflected content likely visible to a new user encountering seed oil discourse. Posts were collected between May 28, 2025, and June 4, 2025, with inclusion limited to content published from January 1, 2020, through the collection period. To capture high-reach discussions, only posts from accounts with more than 5,000 followers were included.

All collected posts were publicly accessible at the time of retrieval, with no private accounts included. Usernames and identifying metadata were removed prior to analysis. Post captions and metadata were extracted using the *Apify* Instagram Post Scraper Actors<sup>1</sup>. The complete dataset, along with annotations and error analysis, is publicly available on Zenodo under a CC BY 4.0 License to support reproducibility and future research<sup>2</sup>.

Our study focuses on caption text rather than image/video content to ensure privacy, feasibility, and reproducibility. Captions can be anonymized and shared with minimal ethical concerns compared to images, which may reveal identifiable individuals or sensitive health information. Instagram captions were preprocessed in Python using *pandas* and *re* library. All scripts were executed in Google Colab (Python 3.12.12) with fixed random

seeds to ensure reproducibility. Preprocessing included: (a) Removing non-alphanumeric characters (excluding basic punctuation), (b) Normalizing whitespace/line breaks, and (c) Preserving emojis as separate tokens.

### 3.2 Human Annotation

We adopt a definition of health misinformation consistent with prior work, defined as “any health-related claim of fact that is false based on current scientific consensus” (Sylvia Chou et al., 2020). In this study, scientific consensus was operationalized according to the 2020–2025 Dietary Guidelines (Dietary Guidelines for Americans, 2020). Posts that contradicted or misrepresented these guidelines were labeled *Misinformation*; posts consistent with or not in conflict with the guidelines were labeled *Credible*.

Three expert annotators with formal training in nutrition independently evaluated each post using a detailed coding manual derived from the U.S. Dietary Guidelines for Americans 2020–2025 (see full annotation guidelines in Appendix A). Annotators were blind to each other’s labels and annotated posts independently in randomized order to minimize order effects. The inter-rater reliability (IRR) among the three annotators was assessed using Fleiss’  $\kappa$  and the final label for each post was determined via majority vote. We treat this human consensus as the operational gold-standard for determining LLM errors. While the annotation pipeline is reproducible at the procedural level, the expert labeling step necessarily involves inter-annotator variability due to subjective judgment.

Figure 2 shows the distribution of caption lengths for misinformation and credible posts. Captions are substantial in length (overall mean: 146–166 words), requiring contextual reasoning beyond isolated claims. Misinformation posts were slightly shorter (M = 142.9, SD = 102.2) than credible posts (M = 176.9, SD = 112.9), though both categories include medium- and long-form captions exceeding

<sup>1</sup><https://apify.com/actors>, last accessed 06/04/2025.

<sup>2</sup><https://doi.org/10.5281/zenodo.20141371>

several hundred words.

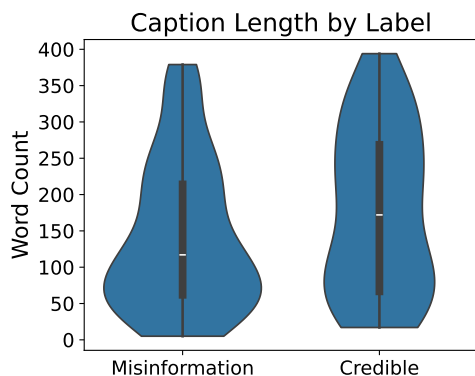


Figure 2: Caption length (word count) by expert label. Both misinformation/ credible posts contain substantial textual content, with mean lengths exceeding 140 words.

Of the 169 posts, 71% (121/169) were labeled as *Misinformation* by expert consensus, indicating class imbalance toward misinformation. All captions were manually reviewed. While some were very short or non-informative, we retained them to reflect real-world variability in Instagram captions.

### 3.3 LLM Annotation

#### 3.3.1 Models

We employed five open-weight, instruction-tuned LLMs: GPT-OSS-13B, Mistral-7B-Instruct, Llama-2-7B, Zephyr-7B, and Qwen3-4B. We focus on open-weight models to ensure transparency and reproducibility of the annotation pipeline. These models span parameter scales from 4B to 13B and were selected to balance instruction-following capability and feasibility under moderate GPU constraints. Default decoding parameters were used; temperature and sampling settings were not manually overridden. Because sampling-based decoding was employed, minor stochastic variation may occur across runs.

#### 3.3.2 Computational Environment

Preprocessing, annotation aggregation, and LLM inference were conducted in Google Colab (Linux-6.6.113+, x86\_64) with approximately 15GB of GPU VRAM. Python version: 3.12.12, pandas version: 2.2.2, NumPy version: 2.0.2, re version: 2.2.1, and Ollama version: 0.17.0 were used.

LLMs were executed locally via the Ollama framework. Inference was conducted sequentially to manage memory constraints. Model outputs were parsed deterministically using regular expressions to extract structured labels and confidence

scores. Inference runtime for annotating the full dataset (169 posts) varied by model size: Qwen3-4B required approximately 30 minutes; Mistral-7B, Llama-2-7B, and Zephyr-7B required 45–60 minutes; GPT-OSS-13B required 90 minutes. All models were executed sequentially.

#### 3.3.3 Prompting Strategy

We experimented with several prompting strategies, including zero-shot, few-shot (1–5 labeled examples of credible and misinformation posts), and knowledge-augmented prompts incorporating brief excerpts from the 2020–2025 U.S. Dietary Guidelines. Across these variants, a hybrid CoT plus knowledge-augmented approach produced the most consistent and interpretable outputs, providing step-by-step reasoning that closely aligned with expert judgment.

Although we explored these alternatives during pilot development, the goal of this study was not prompt optimization but analysis of LLM annotation behavior under a fixed and interpretable configuration. All models were therefore evaluated using the same final prompting setup to ensure comparability across systems.

In our final setup, each model was conditioned as a registered dietitian and instructed to: (1) Generate step-by-step reasoning explaining its label, (2) Output a binary label (0 = credible, 1 = misinformation), and (3) Provide a confidence score, range [0.0, 1.0].

Embedding guideline excerpts in the prompt reduced normative drift and encouraged guideline-grounded reasoning. This hybrid strategy enabled transparent evaluation of LLM annotation behavior and supported systematic error analysis. The full prompt template appears in Appendix B. Outputs were parsed using Python’s `re` library to extract numeric labels and confidence scores. Reasoning text was preserved for qualitative error analysis but did not influence label extraction.

#### 3.3.4 LLM Accuracy Evaluation Protocol

Model label predictions were evaluated against expert consensus using: Accuracy, Precision, Recall, Macro-averaged F1, and Cohen’s  $\kappa$ . Inter-LLM agreement was measured using Fleiss’  $\kappa$ . The LLM–human disagreement rate is computed by comparing the aggregated LLM prediction (majority vote across the five models) against the final expert consensus label for each post.

### 3.4 Error Taxonomy

We developed a hierarchical error taxonomy to characterize misclassifications of Instagram posts about seed oils (full taxonomy is provided in Appendix C). The taxonomy has three levels: (1) error direction, indicating whether a model over- or under-flagged misinformation; (2) error mechanism, describing broad types of model mistakes; and (3) contributing factors, capturing contextual or post-level features that may have influenced the misclassification.

The hierarchical taxonomy captures not only factual correctness but also linguistic and semantic nuances, such as claim hedging, pragmatic interpretation, and subtle contextual cues. These features reflect core challenges in linguistic annotation and provide structured insights into error behavior within the nutrition misinformation domain.

#### 3.4.1 Level 1: Error Direction

Error direction distinguishes between **False Positives (FP)**: The model labeled a credible post as misinformation, and **False Negatives (FN)**: The model labeled a misinformation post as credible. This level captures tendencies toward over-flagging versus permissive behavior, which are relevant for real-world deployment.

#### 3.4.2 Level 2: Error Mechanisms

Each misclassified post was assigned exactly one primary Level-2 error category according to a pre-defined coding manual (see Appendix C). Level-2 categories were defined as mutually exclusive error mechanisms. Annotators followed a hierarchical decision procedure to identify the dominant source of failure. First, errors were evaluated for Content Interpretation Errors (A), defined as cases where the model misread or failed to correctly capture the meaning or implied claims of the post. If the post was correctly interpreted but the model applied incorrect logical, causal, or evaluative reasoning, the error was coded as a Reasoning Error (B). If the reasoning process was coherent but the decision relied on incorrect, missing, or outdated domain-specific nutrition knowledge, it was coded as a Factual Knowledge Error (C). Linguistic and Stylistic Biases (D) (e.g., tone, framing, or presentation effects influencing classification) were included in the taxonomy but were not observed in the present dataset. When multiple mechanisms were present, annotators selected the dominant cause of the final misclassification based on the coding guidelines.

Ambiguous cases were resolved during calibration sessions.

#### 3.4.3 Level 3: Contributing Factors

For each error, the coders also annotated contributing factors at the post level.

**Evidence Type:** Indicates whether and how supporting evidence is provided in the post.

- **NONE:** No sources or supporting evidence.
- **CITED:** External sources such as articles, studies, or links are referenced.
- **DATA:** Quantitative evidence such as statistics, charts, or numerical results are presented.

**Claim Strength:** Captures the degree of certainty expressed in the claim.

- **ABSOLUTE:** Claims expressed with strong certainty or definitive language.
- **QUALIFIED:** Claims framed with tentative or cautious wording.
- **OPINION:** Statements reflecting personal beliefs, preferences, or subjective judgments.

Additional automated characteristics were computed to support exploratory analysis.

#### **Textual Technicality (Flesch Reading Ease):**

The Flesch Reading Ease (FRE) score was calculated using the `textstat` Python package. FRE quantifies readability based on sentence length and word syllables, with higher scores indicating easier-to-read text. Posts were categorized into three levels based on the FRE score:

- **LAY** ( $\geq 60$ ): Texts that are simple and easily understood by a general audience.
- **MIXED** (40–59): Moderate complexity.
- **TECHNICAL** ( $< 40$ ): Texts that are dense, highly technical, or difficult to read.

**Post Length:** Measured in characters using Python, excluding spaces, punctuation, and special characters. Posts were categorized into:

- **SHORT** ( $< 100$  chars): Very brief posts.
- **MEDIUM** (100–500 chars): Moderate length.
- **LONG** ( $> 500$  chars): Extended posts.

#### 3.4.4 LLM Errors Coding Procedure

Two coders independently annotated a 20% random subset of misclassified posts. Inter-rater reliability was computed using Cohen’s  $\kappa$ . Coders then

met to resolve disagreements and clarify category definitions, with an expert guiding calibration to ensure consistent interpretation. The remaining errors were annotated by a single trained coder using the agreed-upon taxonomy, revisiting ambiguous cases as needed. The full coding manual is provided in Appendix C.

### 3.4.5 Error Multi-Factor Analysis

We first used Level 1 to identify all misclassified predictions. Level 2 then characterized the type of error for each model (content interpretation, reasoning, factual knowledge), while Level 3 examined post-level features: claim strength, evidence type, textual style, and caption length for each misclassification. Level-3 characteristics were counted once per misclassified prediction. Combining Level-2 error mechanisms with Level-3 features allowed us to explore how content and surface level factors relate to model specific error tendencies, including FP and FN. Confidence-weighted analyses were conducted to examine how LLM self-reported model confidence correlates with error direction.

## 4 Results

### 4.1 Human Annotation Reliability

Three nutrition-trained annotators independently labeled 169 Instagram posts. Pairwise Cohen’s  $\kappa$  ranged from 0.789 to 0.828, with Fleiss’  $\kappa = 0.813$ , indicating substantial agreement. The human disagreement rate (11.8%) is computed as the proportion of instances on which at least one of the three expert annotators disagreed prior to majority voting. This reflects a set of borderline cases with inherent annotation ambiguity. We include a representative case study of annotator disagreement in Appendix D, illustrating how such borderline claims map onto our hierarchical error taxonomy.

### 4.2 LLM Annotation Performance

#### 4.2.1 Inter-Model Agreement

Inter-model agreement among the five LLMs was lower than human agreement, with pairwise Cohen’s  $\kappa$  spanning 0.237 to 0.813 and Fleiss’  $\kappa = 0.417$ , suggesting moderate consistency. While over half of all LLM predictions (56.8%) disagreed with at least one other LLM, only 11.8% of predictions differed from the human consensus labels. As a descriptive observation, the identical magnitude of these two rates (11.8%) is notable in this dataset, although it arises from different underlying pro-

cesses. This LLM–human alignment rate is comparable in magnitude to the human inter-annotator disagreement rate reported in Section 4.1, although the two measures reflect different evaluation procedures and are not directly equivalent.

This pattern indicates that LLMs generally align with human judgment on clear-cut cases, with most disagreements arising from inherently ambiguous posts. As Figure 3 shows, full agreement across all five LLMs occurred for 42% of posts, with multi-model disagreement being rare.

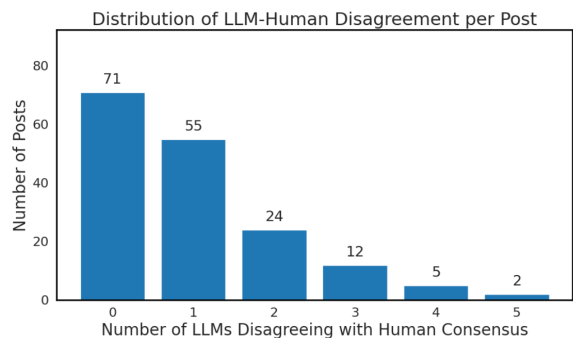


Figure 3: Distribution of the number of LLMs disagreeing with human consensus.

#### 4.2.2 Individual LLM Performance

Table 1 reports individual model metrics against human consensus. GPT-OSS and Qwen3 achieved the highest F1 and Cohen’s  $\kappa$ , whereas Llama-2 performed substantially worse than the other models, with notably lower F1 score and Cohen’s  $\kappa$ , indicating weaker agreement with human annotations. Zephyr and Mistral showed intermediate performance, with Mistral achieving relatively balanced metrics across measures.

Model	Acc.	Prec.	Rec.	F1	$\kappa$
GPT-OSS	0.905	0.973	0.893	0.931	0.781
Mistral	0.852	0.864	0.942	0.901	0.609
Llama-2	0.615	0.811	0.603	0.692	0.209
Zephyr	0.746	0.933	0.694	0.796	0.476
Qwen3	0.882	0.939	0.893	0.915	0.720

Table 1: Individual LLM performance against human consensus.  $\kappa$  is Cohen’s  $\kappa$ .

### 4.3 LLM Error Analysis

Of the 169 Instagram posts, 98 contained at least one LLM prediction that disagreed with the expert consensus. These 98 posts account for 169 of the 845 total predictions (169 posts  $\times$  5 models, 20%). For Level-2 analysis, we examine each

misclassified prediction using our hierarchical error taxonomy to systematically characterize model failure modes. Level-3 analysis considers post-level features across these 98 posts to identify content characteristics associated with LLM errors.

### 4.3.1 Level 1 Error Direction

Across all models, we observed 169 misclassifications, with substantially more FN ( $n = 118$ ) than FP ( $n = 51$ ), indicating a general tendency to under-flag misinformation. Model-level comparison shows distinct behavioral tendencies: models such as Llama-2 and Zephyr produced a higher proportion of FNs, whereas GPT-OSS and Mistral exhibited a more balanced FP/FN distribution. Figure 4 illustrates the distribution of FN and FP across models.

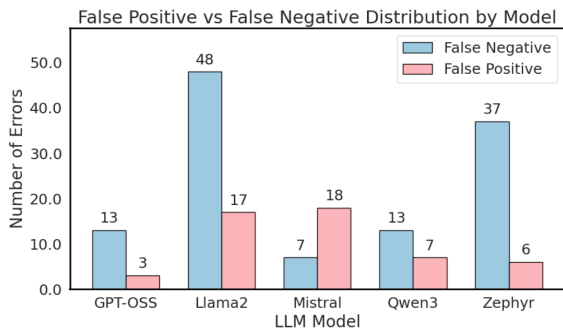


Figure 4: False positive vs. false negative distribution.

### 4.3.2 Inter-Coder Reliability for Errors

For a subset of misclassified posts (20%) independently annotated by two coders, inter-rater agreement was substantial, with Cohen’s  $\kappa = 0.774$  for primary error mechanism (Level 2) codes. For post-level contributing factors, agreement was  $\kappa = 0.726$  for Evidence Type and  $\kappa = 0.758$  for Claim Strength. Disagreements were resolved through discussion with an expert coder. Following this calibration, the second coder completed annotation of the remaining posts using the clarified taxonomy, ensuring consistent coding across the full dataset.

### 4.3.3 Level 2 Error Mechanisms

Content interpretation errors dominated LLM annotation failures, followed by factual knowledge and reasoning errors (see Figure 5). No errors were attributed to linguistic or stylistic bias in this dataset; however, our hierarchical taxonomy is designed to capture such issues in corpora where style influences annotation, making it broadly applicable.

As Figure 5 shows, most failures arise from challenges in understanding the meaning, nuance, and implied claims in social media posts, highlighting areas where LLMs require additional grounding or contextual reasoning to align with expert judgment.

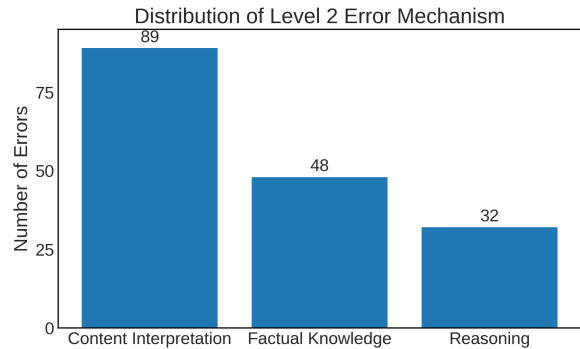


Figure 5: Distribution of Level 2 error mechanisms across all misclassified LLM predictions.

### 4.3.4 Level 3 Post Characteristics

Analysis of post-level features (Figure 6) shows that most posts contained QUALIFIED claims, with fewer ABSOLUTE statements or OPINION-based content. Evidence patterns show that the majority of posts contained NONE evidence, while only a small fraction included DATA or CITED sources. Textual style was primarily LAY or MIXED, with fewer TECHNICAL posts. Most captions were relatively LONG. This overview highlights the typical contexts in which LLM annotation was applied, providing insight into the types of content that pose challenges for automated labeling. All counts reported here are based on the 98 posts that contained at least one LLM prediction that disagreed with the expert consensus.

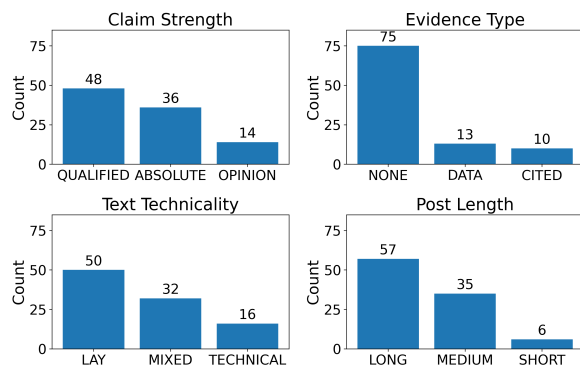


Figure 6: Distribution of Level 3 post characteristics across the dataset, including claim strength, evidence type, textual technicality, and caption length.

#### 4.4 Multi-Factor Error Analysis

Across models, the largest number of misclassifications occurred for Llama-2 ( $n = 65$ ), followed by Zephyr ( $n = 43$ ) and Mistral ( $n = 25$ ), while Qwen3 ( $n = 20$ ) and GPT-OSS ( $n = 16$ ) produced fewer errors overall. Figure 7 illustrates the breakdown of misclassifications by Level-2 error type for each model.



Figure 7: Breakdown of misclassifications by Level 2 error type for each model. Content interpretation errors (A), reasoning errors (B), and factual knowledge errors (C) are shown.

Analysis of Level 3 post characteristics reveals systematic associations with specific Level 2 errors. Level 3 characteristics are counted once per misclassified prediction; thus, posts with multiple misclassifications contribute multiple observations. This results in 169 misclassifications derived from 98 posts (see Appendix E for detailed heatmaps).

Content Interpretation Errors (A) were most common in posts containing ambiguous or nuanced claims. Error rates increased when claims used tentative language (QUALIFIED) or lacked supporting evidence (NONE), suggesting that models struggle with hedged or context-dependent statements. Posts written in LAY language and those with LONG captions were also more frequently misinterpreted, indicating that discourse complexity may contribute to annotation errors.

Reasoning Errors (B) occurred primarily in posts relying on anecdotal or loosely causal reasoning. These errors were especially common in posts containing QUALIFIED or OPINION-based claims, suggesting difficulty evaluating tentative or context-dependent arguments. Factual Knowledge Errors (C) were most often associated with posts requiring domain-specific nutrition knowledge. Posts lacking supporting evidence (NONE) accounted for most factual knowledge failures, whereas posts contain-

ing CITED or DATA sources were more reliably annotated. This pattern suggests that explicit evidentiary cues help anchor model judgments and reduce knowledge-related annotation errors.

Across error types, post-level characteristics consistently influenced misclassification patterns. QUALIFIED claims produced the largest number of errors overall, while ABSOLUTE claims more frequently resulted in FN errors. The absence of supporting evidence exacerbated all error types. LONG captions and LAY style language were also associated with higher error rates, likely reflecting increased discourse complexity. These results indicate that LLM annotation errors are shaped by interactions between linguistic framing, evidential context, and domain knowledge requirements.

#### 4.5 Confidence Calibration

LLM self-reported confidence varied across models. Calibration curves in Figure 8 illustrate that GPT-OSS and Zephyr are relatively well-aligned with empirical correctness, whereas Llama-2 and Mistral exhibit over- or under-confidence. Notably, of 169 total misclassified outputs, 149 were reported with confidence  $\geq 0.8$ , indicating systematic overconfidence.

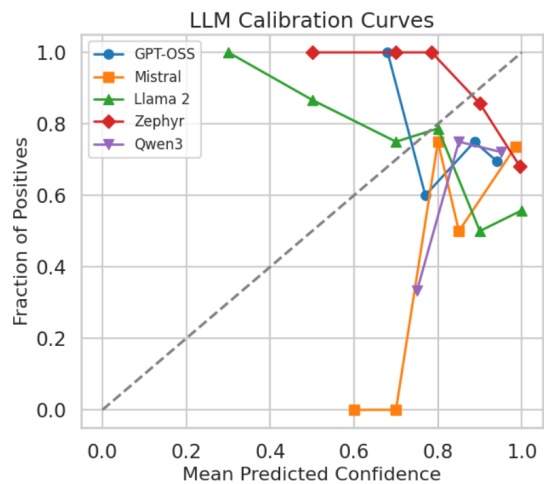


Figure 8: Calibration curves comparing predicted LLM confidence vs. empirical correctness.

## 5 Conclusion

We present a reproducible, expert-anchored workflow for annotating health misinformation on social media using LLMs, complemented by a hierarchical error taxonomy that categorizes annotation failures into interpretable mechanisms. Our results show that LLM errors are systematic rather than

random, arising primarily from miscalibration of claim strength, insufficient grounding in evidence, and challenges in interpreting hedged or absolute language. Persistent overconfidence in incorrect predictions indicates that failures are driven by pragmatic interpretation and epistemic misalignment, rather than surface-level textual complexity. The taxonomy captures linguistic and semantic nuances, such as claim hedging and subtle contextual cues, highlighting challenges that extend beyond nutrition misinformation and informing general best practices in linguistic annotation.

By linking error mechanisms to post characteristics and model behavior, this workflow provides guidance for when LLMs are reliable and for best practices in LLM-based annotation, including knowledge-augmented prompting and step-by-step reasoning aligned with expert judgment. It also supports transparent analysis of human-machine disagreement, enabling systematic evaluation of annotation quality and illuminating recurring failure modes. Overall, this study provides a reproducible and linguistically grounded insight for evaluating LLM annotations, offering insight for improving annotation reliability in domain-specific NLP tasks.

Future work should scale the workflow to larger and more diverse corpora, explore multimodal reasoning, and validate taxonomy coding with multiple annotators to strengthen reliability and generalizability. Future work may also incorporate active-learning loops in which LLM explanations trigger selective re-annotation, and evaluate ensemble disagreement as an uncertainty signal. These directions would deepen understanding of both annotation reliability and model-driven quality control strategies in domain-specific contexts.

## Limitations

This study has several limitations. The dataset is modest in size (169 posts), which may limit the generalizability of the quantitative findings. The study focuses on expert-driven annotation and qualitative error analysis within a specific domain of nutrition misinformation rather than large-scale model evaluation. Additionally, the dataset is topic-specific (“seed oils”), English-only, and grounded in U.S. dietary guidelines, which further limits generalizability across domains, cultures, and platforms. Posts were restricted to accounts with more than 5,000 followers to prioritize high-reach content that is more likely to influence public discourse. How-

ever, this restriction may introduce sampling bias and may not fully capture misinformation circulating in smaller or private communities.

The annotation protocol relies on three expert annotators and majority voting to determine gold labels. Although inter-rater reliability was substantial (Fleiss’  $\kappa = 0.813$ ), the disagreement rate indicates that some posts are inherently ambiguous, and the gold labels should be interpreted as consensus judgments rather than absolute truth.

Although binary labeling aligns with prior misinformation detection work and enables consistent evaluation, health-related claims often exist on a spectrum rather than as strictly true or false categories. Posts that combine accurate information with misleading interpretations may be especially difficult to annotate. Future work could explore ordinal or multi-label annotation schemes to better capture such nuance and annotator uncertainty.

The study evaluates a limited set of five open-weight language models under a single prompting configuration. Model outputs may vary with different prompts, decoding parameters, or repeated runs due to stochasticity, and the analysis does not include larger proprietary or multimodal models, which may behave differently. The error taxonomy involves interpretive coding, and although inter-coder reliability was measured and calibration performed, some subjectivity is inherent in categorizing errors.

Future work may extend this evaluation to larger proprietary or multimodal models to examine whether similar annotation error patterns emerge across different model architectures. Finally, the analysis focuses exclusively on caption text, whereas Instagram is inherently multimodal; visual content may influence how claims are interpreted, so these findings pertain specifically to textual linguistic annotation.

## Ethical Considerations

This study uses only publicly accessible Instagram posts and complies with platform policies and institutional research standards. No private accounts were accessed, and all identifying information was removed prior to analysis. Any released dataset will contain anonymized caption text only.

We acknowledge that misinformation detection systems may produce false positives or false negatives with uneven social consequences. The models evaluated here are intended for research purposes

and should not be deployed for fully automated moderation without human oversight. We discourage applications involving surveillance, censorship, or suppression of legitimate expression.

Our goal is to advance transparency and reproducibility in health misinformation annotation research while prioritizing privacy, fairness, and responsible use.

## Acknowledgments

We thank Charlotte Martin, Asal Abbaszadeh, Richard Loftis, and Alan Flanagan for their valuable contributions to extracting and annotating the original posts. The authors also gratefully acknowledge the School of Computing, College of Computing, College of Computing, Engineering and Construction, and Graduate School at University of North Florida for their support and funding, which made this research possible.

## References

- Meysam Alizadeh, Maël Kubli, Zeynab Samei, Shirin Dehghani, Mohammadmasiha Zahedivafa, Juan D Bermeo, Maria Korobeynikova, and Fabrizio Gilardi. 2025. Open-source llms for text annotation: a practical guide for model setting and fine-tuning. *Journal of Computational Social Science*, 8(1):17.
- Joseph Balagtas and Elijah Bryant. 2025. *Consumer food insights*. Technical report, Center for Food Demand Analysis and Sustainability, Purdue University.
- Brian Dean. 2025. *Instagram demographic statistics: How many people use instagram in 2024?* Accessed: 2025-04-25.
- Dietary Guidelines for Americans. 2020. Dietary guidelines for americans, 2020-2025 and online materials. <https://www.dietaryguidelines.gov/resources/2020-2025-dietary-guidelines-online-materials>.
- Tanjim Bin Faruk. 2024. Evaluating the performance of large language models in scientific claim detection and classification. <https://arxiv.org/abs/2412.16486>.
- Akshay Goel, Almog Gueta, Omry Gilon, Chang Liu, Sofia Erell, Lan Huong Nguyen, Xiaohong Hao, Bolous Jaber, Shashir Reddy, Rupesh Kartha, and 1 others. 2023. Llms accelerate annotation for medical information extraction. In *machine learning for health (MLAH)*, pages 82–100. PMLR.
- Prajwol Lamichhane, Indika Kahanda, Andrea Arikawa, Charlotte Martin, Maribel Garcia, Camila Figueiredo, and Haivan Benjamin. 2025. Exploring the feasibility of identifying nutrition misinformation on social media. In *Proceedings of the ACM/IEEE International Conference on Connected Health: Applications, Systems and Engineering Technologies*, pages 319–323.
- Omer Nahum, Nitay Calderon, Orgad Keller, Idan Szpektor, and Roi Reichart. 2025. Are llms better than reported? detecting label errors and mitigating their effect on model performance. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 26770–26797.
- Kristina S Petersen, Kevin C Maki, Philip C Calder, Martha A Belury, Mark Messina, Carol F Kirkpatrick, and William S Harris. 2024. Perspective on the health effects of unsaturated fatty acids and commonly consumed plant oils high in unsaturated fat. *British Journal of Nutrition*, pages 1–12.
- Jiaxing Qiu, Dongliang Guo, Natalie Papini, Noelle Peace, Hannah F Fitterman-Harris, Cheri A Levinson, Tom Hartvigsen, and Teague R Henry. 2025. Labeling free-text data using language model ensembles. *arXiv preprint arXiv:2501.08413*.
- Hope Schroeder, Deb Roy, and Jad Kabbara. 2025. Just put a human in the loop? investigating llm-assisted annotation for subjective tasks. In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 25771–25795.
- Sergio Segado-Fernández, Beatriz Jiménez-Gómez, Pedro Jesús Jiménez-Hidalgo, María del Carmen Lozano-Estevan, and Iván Herrera-Peco. 2025. Disinformation about diet and nutrition on social networks: a review of the literature. *Nutricion hospitalaria*, 42(2).
- Wen-Ying Sylvia Chou, Anna Gaysynsky, and Joseph N Cappella. 2020. Where we go from here: health misinformation on social media.
- Dongmei Tan, Yi Huang, Ming Liu, Ziyu Li, Xiaoqian Wu, and Cheng Huang. 2025. Identification of online health information using large pretrained language models: Mixed methods study. *Journal of Medical Internet Research*, 27:e70733.
- Zhen Tan, Dawei Li, Song Wang, Alimohammad Beigi, Bohan Jiang, Amrita Bhattacharjee, Mansoor Karami, Jundong Li, Lu Cheng, and Huan Liu. 2024. Large language models for data annotation and synthesis: A survey. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 930–957.
- Xinru Wang, Hannah Kim, Sajjadur Rahman, Kushan Mitra, and Zhengjie Miao. 2024. Human-llm collaborative annotation through effective verification of llm labels. In *Proceedings of the 2024 CHI conference on human factors in computing systems*, pages 1–21.
- Andy Wai Kan Yeung, Anela Tosevska, Elisabeth Klager, Fabian Eibensteiner, Christos Tsagkaris, Emil D Parvanov, Faisal A Nawaz, Sabine Völkl-Kernstock, Eva Schaden, Maria Kletecka-Pulker, and 1 others. 2022. Medical and health-related misinformation on social media: bibliometric study of the scientific literature. *Journal of Medical Internet Research*, 24(1):e28152.

## Appendix

### A Annotation Guidelines

#### A.1 Purpose of These Guidelines

These guidelines ensure that all human coders classify social media nutrition posts consistently using:

- A two-level accuracy scale (0 or 1)
- Rules aligned with the *Dietary Guidelines for Americans, 2020–2025 (DGA)* and the *2020 DGAC Scientific Report*

Coders should reference only the content presented in the post and apply the rules exactly as described.

#### A.2 What Coders Evaluate

For each post, coders assign:

1. Accuracy classification (1 or 0)
2. An optional brief justification (2–3 sentences)

Coders do **not** evaluate:

- Creator’s biography
- Hashtags
- Engagement metrics
- Intent or motivation
- User comments
- Information not shown in the post itself

#### A.3 Evidence Standard

Coders must evaluate claims solely against:

- *Dietary Guidelines for Americans 2020–2025*
- *DGAC Scientific Report (2020)*

Coders should **not** use:

- Personal knowledge
- Other nutrition guidelines
- Mechanistic physiology not discussed in the DGA/DGAC
- External sources, papers, or reputable organizations
- Internet searches

#### A.4 The Two-Level Accuracy Classification

**1 - Accurate or Mostly Accurate** A post is **Accurate** when:

- All nutrition claims fully align with DGA/DGAC
- Any minor inaccuracies do not change the overall message

#### Examples of Accurate Content

- Saying omega-6 fats are not inflammatory
- Correcting misinformation about seed oils

A post is **Mostly Accurate** when:

- It is evidence-aligned overall
- It contains some non-trivial inaccuracies, but the main message remains correct

#### Typical patterns

- Slight exaggeration without reversing the evidence

**0 - Mostly Inaccurate or Inaccurate** A post is **Mostly Inaccurate** when:

- It contains some correct information
- Misleading or incorrect claims dominate

#### Typical patterns

- Mixing evidence-based advice with unsupported mechanistic claims
- Overstating inflammation, hormone effects, oxidation, or fear-based messages about oils
- Encouraging elimination of seed oils while acknowledging they contain essential fats

A post is **Inaccurate** when:

- The core message contradicts DGA/DGAC
- Accurate statements, if present, are trivial

#### Examples

- “Butter is heart healthy; seed oils cause disease.”
- “Seed oils are toxic poisons that cause inflammation.”

#### A.5 Reasoning Requirements (2–3 Sentences)

Coders may provide a short justification for doubtful annotations. The justification must:

- Evaluate the overall message rather than line-by-line details
- Reference only areas covered by the DGA/DGAC
- Not mention excluded content (e.g., “the DGA doesn’t cover . . .”)
- Avoid SEO phrases or quoting the creator directly
- Clearly explain why the classification was chosen

#### A.6 Special Situations

**Mixed-accuracy posts** Code based on the **overall message**, not on isolated statements.

**Mechanistic claims (inflammation, oxidation, hormones)** If the mechanism is not discussed in the DGA:

- Ignore the mechanism itself
- Evaluate whether the dietary conclusion contradicts DGA guidance

**Reactive content** If a post stitches another creator's video:

- Ignore the introductory clip
- Code only the creator's commentary

## B Full Prompt Template

### Prompt Template Used for All Models

You are a registered dietitian and an expert in evaluating health-related claims. Can you identify whether the following Instagram post contains misinformation that contradicts the Dietary Guidelines for Americans 2020–2025?

#### Key points from the guidelines:

- Oils provide essential fatty acids and are part of a healthy diet.
- Tropical oils like coconut and palm oils are high in saturated fat.
- Replace high saturated fat foods with vegetable oils.

Post: {post}

Provide step-by-step reasoning.

End with:

Label: 0 if credible, 1 if misinformation

Confidence: 0.0 to 1.0

## C Error Taxonomy Coding Manual

### C.1 Instructions for Coders

This manual guides the classification of errors made by LLMs when annotating Instagram posts about seed oils as *misinformation* or *credible*.

Before coding:

- Familiarize yourself with all error categories and examples.
- Each error should receive **one primary error type**.
- Up to **two secondary contributing categories** may be assigned when appropriate.
- When uncertainty arises, discuss the case with the other coder.

### C.2 Coding Process

1. Read the Instagram post carefully.
2. Review the **ground truth label** determined through human consensus.
3. Note the **LLM's predicted label**.
4. Identify why the LLM made the incorrect classification.
5. Assign the appropriate error category and any contributing factors.

### C.3 Level 1: Error Direction

Errors are first coded by direction:

- **FP (False Positive)**: The LLM labeled the post as misinformation, but the post is credible.
- **FN (False Negative)**: The LLM labeled the post as credible, but the post contains misinformation.

### C.4 Level 2: Error Mechanism Categories

#### A. Content Interpretation Errors

These errors occur when the LLM misinterprets the meaning or intent of the content.

**Nuance Blindness**: Failure to detect linguistic cues such as sarcasm, irony, or hedging that alter the meaning of the post.

- “Yeah sure, seed oils are definitely going to kill us all ”
- “Love how everyone became a biochemist overnight to explain seed oils”

**Context Insensitivity**: Failure to incorporate contextual signals that influence interpretation.

- Ignoring hashtags such as #sarcasm or #justkidding.
- Treating posts from expert accounts the same as unverified sources.

**Claim Granularity Errors**: Failure to distinguish between neutral factual statements, nuanced claims, and extreme assertions.

- “Seed oils contain omega-6 fatty acids” labeled as misinformation.
- “I avoid seed oils” treated as a factual health claim rather than a personal preference.

**Source Credibility Misjudgment**: Incorrect evaluation of cited sources or reliance on presentation style rather than evidence quality.

- Credible health organization quoted but ignored by the model.
- Misinformation written in academic style treated as credible.

## B. Reasoning Errors

These errors arise from flawed reasoning processes when evaluating claims.

**Over-Generalization:** Applying overly broad patterns instead of evaluating the specific claim.

- Automatically labeling any criticism of seed oils as misinformation.
- Pattern matching: “seed oils” + “inflammatory” leading to automatic misclassification.

**Under-Generalization:** Accepting problematic claims as credible because they are framed cautiously.

- “In my experience seed oils caused my health problems” treated as credible evidence.
- “I’m just asking questions about seed oils” interpreted as neutral inquiry.

**Logical Fallacy Blindness:** Failure to detect flawed reasoning such as anecdotal evidence or correlation–causation confusion.

- “I stopped eating seed oils and my acne disappeared, therefore seed oils cause acne.”
- Anecdotal claims treated as scientific evidence.

## C. Factual Knowledge Errors

These occur when the LLM lacks correct domain knowledge or invents information.

**Outdated Knowledge:** Reliance on obsolete information or outdated scientific consensus.

- Referencing outdated dietary guidance rather than more recent recommendations.

**Factual Hallucination:** Confidently asserting incorrect facts or nonexistent evidence.

- Claiming a specific study proves seed oils are toxic when no such study exists.

**Domain Knowledge Gaps:** Insufficient understanding of nutrition or biochemical concepts required to evaluate claims.

- Confusion between omega-3 and omega-6 fatty acids.
- Misinterpreting oxidative stress mechanisms.

## D. Linguistic and Stylistic Biases

These errors occur when stylistic features influence classification rather than factual accuracy.

**Formality Bias:** Judgments influenced by writing style instead of claim validity.

- Academic-sounding misinformation labeled as credible.
- Accurate information written casually labeled as misinformation.

**Confidence Confusion:** Confusing the author’s confidence level with factual correctness.

- Highly confident misinformation treated as credible.
- Cautious scientific language interpreted as uncertainty or misinformation.

**Emotional Language Sensitivity:** Overreacting to emotional tone rather than evaluating factual content.

- Fearful language triggering a misinformation label even when the claim is accurate.

## C.5 Contributing Factors

For each error, coders may optionally record contextual factors that help explain the misclassification.

### • Evidence Provided

- **NONE:** No sources or supporting evidence are provided.
- **CITED:** External sources (e.g., articles, studies, organizations) are referenced.
- **DATA:** Quantitative evidence such as statistics, charts, or graphs is presented.

### • Claim Strength

- **ABSOLUTE:** Strong certainty (e.g., “always”, “never”, “proves”).
- **QUALIFIED:** Tentative or cautious wording (e.g., “may”, “might”, “suggests”).
- **OPINION:** Personal beliefs or preferences (e.g., “I think”, “in my experience”).

## C.6 Inter-Rater Reliability Protocol

To ensure consistent application of the hierarchical error taxonomy, the following procedure should be followed:

1. Two coders independently annotate a randomly selected **20% subset** of the dataset using the hierarchical error taxonomy described above.
2. After both coders complete their annotations, inter-rater reliability should be calculated using **Cohen’s  $\kappa$** .
3. The coders then meet to review cases of disagreement and discuss the reasoning behind their coding decisions.
4. During this discussion, an expert coder clarifies category definitions and resolves ambiguities in the taxonomy to ensure that both coders interpret the categories consistently.
5. After this calibration discussion, the remaining errors in the dataset are annotated by the trained coder using the agreed-upon interpretations of the taxonomy.

### C.7 Illustrative Examples

#### Example 1

**Post:** “So apparently seed oils are the reason I’m tired all the time? Sure Jan”

**Ground Truth:** Credible

**LLM Label:** Misinformation

**Error Direction:** FP

**Error Mechanism:** A. Content Interpretation Errors

**Explanation:** LLM missed sarcasm (“Sure Jan” meme reference)

**Contributing Factors:**

- Length: SHORT
- Technicality: LAY
- Evidence: NONE
- Claim Strength: ABSOLUTE

#### Example 2

**Post:** “I stopped using seed oils 6 months ago and my inflammation completely disappeared. Do your own research!”

**Ground Truth:** Misinformation

**LLM Label:** Credible

**Error Direction:** FN

**Error Mechanism:** B. Reasoning Errors

**Explanation:** LLM failed to identify correlation/causation fallacy and accepted anecdotal evidence as proof

**Contributing Factors:**

- Length: MEDIUM
- Technicality: MIXED
- Evidence: ANECDOTAL
- Claim Strength: ABSOLUTE

## D Human Disagreement Case Study

To illustrate sources of annotation ambiguity and how they relate to our hierarchical error taxonomy, we present a representative case from the subset of posts where expert annotators initially disagreed prior to adjudication.

### D.1 Post Example

“Plot twist not all oils are created equal! Say goodbye to those heavily processed seed oils that are lurking in your pantry. Your body (and your tastebuds) will thank you for choosing whole food alternatives like olive oil coconut oil and avocado oil instead instead. Who else is joining the real food revolution? Real-Food HealthyChoices NutritionTips SeedOils HealthySwaps WellnessJourney”

### D.2 Human Annotation Disagreement

This post was independently annotated by three nutrition experts. One annotator labeled the post as **Credible**, interpreting the content as a general dietary suggestion promoting whole food alternatives without making an explicit false factual claim. In contrast, two annotators labeled the post as **Misinformation**, focusing on the implied negative framing of seed oils as “heavily processed” and the absence of scientific evidence supporting the suggested health inference.

Following majority voting, the final gold label for this instance is **Misinformation**.

### D.3 LLM Prediction Error

Among the evaluated models, only Llama-2 incorrectly classified this post as **Credible**, constituting a False Negative (FN) with respect to the human consensus label.

### D.4 Taxonomy-Based Analysis

Using our hierarchical error taxonomy, this case is categorized as follows:

**Level 1 (Error Direction):** False Negative (FN)

**Level 2 (Error Mechanism):** Content Interpretation Error

The misclassification arises from divergent interpretation of the post’s implied claims. While the

model appears to focus on the absence of explicit factual assertions, human annotators considered the implicit causal framing (i.e., seed oils as harmful) as sufficient for a misinformation label.

**Level 3 (Contributing Factors):**

- **Claim Strength:** QUALIFIED / IMPLIED  
ABSOLUTE
- **Evidence Type:** NONE
- **Textual Technicality:** LAY
- **Post Length:** Short

## **D.5 Discussion**

This example highlights a common source of disagreement in nutrition-related misinformation annotation: posts that combine factual product substitutions with implicit health claims. Although no explicit scientific claim is made, the framing of seed oils as “heavily processed” and the endorsement of alternative oils introduces an evaluative inference that leads to divergence in human interpretation. This case illustrates how Content Interpretation Errors in LLMs often stem from failures to capture pragmatic implications rather than explicit factual inaccuracies.

## **E Level 2 and Level 3 Error Heatmaps**



Figure 9: Heatmaps showing associations between Level 2 error mechanisms (content interpretation, reasoning, factual knowledge) and Level 3 post-level features (claim strength, evidence type, textual style, caption length) across misclassified predictions. Each panel represents a separate post-level feature.

# Math-DB: A Discourse Framework for Mathematical Word Problems to Enhance LLM Reasoning

Mustafa Erolcan Er

Department of Cognitive Science  
Middle East Technical University (METU)  
Ankara, Türkiye  
erolcan@metu.edu.tr

## Abstract

Large Language Models have demonstrated significant progress in solving mathematical word problems through techniques like Chain-of-Thought (CoT) prompting. However, recent research indicates that these models often rely on statistical regularities and surface-level patterns rather than true logical reasoning, leading to performance drops when faced with minor problem perturbations or irrelevant information. In this study, we introduce Math Discourse Bank (Math-DB), a novel discourse framework and annotated dataset designed to enhance LLM reasoning. Inspired by the Penn Discourse TreeBank (PDTB) and mathematics education research, Math-DB defines a hierarchy of discourse senses designed for quantitative reasoning, including categories such as Change, Combine, Compare, and Equalize. We applied this framework to the GSM-Symbolic dataset of 12,500 problems, yielding 47,815 sense-labeled discourse relations over 11,414 successfully-aligned instances (91.3% pipeline yield). Our experiments demonstrate that incorporating Math-DB annotations into CoT prompts consistently improves LLM performance across various difficulty levels.

## 1 Introduction

Large Language Models have made significant improvement in solving mathematical word problems, especially with techniques like CoT prompting (Wei et al., 2022). In particular, prompting LLMs to generate intermediate reasoning steps has enabled remarkable performance on benchmarks such as GSM8K (Cobbe et al., 2021). Despite these advances, a growing body of evidence indicates that current LLMs still fall short of true logical reasoning. By *true logical reasoning*, we mean the ability to derive a correct solution from the underlying semantic structure of a problem in a way that is invariant to surface-level changes such as renaming entities, substituting numerical values, or inserting

irrelevant context. Empirically, this is distinguished from template- or pattern-based behavior by testing whether model accuracy remains stable under such controlled perturbations. Current LLMs instead often detect statistical regularities and template patterns present in the training data. This can lead to memorized behavior, where small changes in a problem’s presentation can cause drastic drops in LLM’s performance (Mirzadeh et al., 2024).

Mirzadeh et al. (2024) generated diverse math problems from symbolic templates to test whether models truly understand the reasoning or just the surface form. They found that even minor perturbations, such as modifying numerical values or proper nouns, noticeably degrade LLM performance. Furthermore, inserting irrelevant information (distractor clauses that do not affect the solution) can confuse models into drastic errors. This suggests that models do not effectively reason about math word problems but rather memorize the answers already found in training data. Similar findings by Shi et al. (2023) were obtained on GSM-IC, a variant of GSM8K in which each problem is augmented with a sentence containing irrelevant but topically related information; the addition of such distractor context caused dramatic performance decreases.

One promising direction is to provide semantic guidance for the LLM’s reasoning. Humans solving word problems rely on the narrative relations between quantities (e.g. identifying that one quantity increases another, or that two parts combine into a whole) rather than just crunching numbers in order. Discourse frameworks like PDTB (Prasad et al., 2017) capture how clauses connect logically (causal, temporal, contrastive), but general-purpose relations such as *Contingency.Cause* do not directly map to the operations in math problems. Mathematics education research, by contrast, has long categorized arithmetic word problems into schema types based on the semantic relations between quantities (Daroczy et al., 2015; Riley et al., 1984). We

adopt and adapt these categories in Math-DB (Section 3), hypothesizing that labeling the semantic relations between clauses can guide the model to focus on what each step means rather than relying on positional heuristics or keywords alone.

In this paper, we present Math-DB, a novel discourse framework and annotated dataset for mathematical word problems. Math-DB defines a PDTB-style hierarchy of discourse senses customized for quantitative reasoning. Each pair of adjacent clauses in a math problem is assigned a relation sense that describes its math-relevant semantic connection. A sample annotated problem can be seen in Figure 1 below.

We applied Math-DB to the full GSM-Symbolic benchmark of Mirzadeh et al. (2024), comprising 5,000 main problems, 5,000 GSM-P1 problems, and 2,500 GSM-P2 problems, for a total of 12,500 source instances. GSM-P1 problems are generated by adding one additional clause to each original problem, while GSM-P2 problems add two such clauses. Each added clause introduces one extra reasoning step, forcing the model to perform additional calculations beyond the standard baseline question. Our semi-automatic annotation pipeline successfully aligned 11,414 of these 12,500 instances (4,621 main, 4,535 P1, and 2,258 P2; 91.3% overall yield), producing 47,815 sense-labeled discourse relations. The remaining 1,086 instances were rejected at the alignment stage and are documented in the released diagnostic logs (see §4). To create Math-DB, we built a semi-automatic pipeline: a fine-tuned BERT-based parser identifies discourse connectives and segments each problem into PDTB-style arguments (Arg-1, Arg-2), after which six trained undergraduate annotators from a mathematics department assign the relation sense to each link.

Finally, we show how Math-DB can enhance Chain-of-Thought prompts: alongside the problem text, the model receives discourse annotations that flag the operation each clause introduces (e.g. aggregating parts vs. subtracting a removed amount). Our results (Section 5.2) show that this guidance makes the model more robust to the added-complexity perturbations in GSM-P1 and GSM-P2.

In summary, our contributions are:

- A domain-specific discourse relation hierarchy for math word problems, grounded in known semantic problem types (Change,

Combine, Compare, Equalize, etc.), presented in Section 3.

- An annotated corpus consisting of 47,815 Math-DB discourse relations over 11,414 successfully-aligned instances from the 12,500-problem GSM-Symbolic benchmark, together with diagnostic logs covering the 1,086 unaligned instances, described in Section 4. This new corpus will be released publicly, providing a resource to train and evaluate models that can understand and predict mathematical discourse structure.
- A discourse-augmented prompting experiment that elaborates classical CoT prompts with Math-DB annotations for LLMs, presented in Section 5.

## 2 Related Work

### 2.1 Discourse Frameworks

Clause connectivity is a fundamental focus in NLP, exemplified by the Penn Discourse Treebank (PDTB) (Prasad et al., 2017). The PDTB framework utilizes a lexically grounded, three-level hierarchy to map logical and pragmatic links between abstract objects (Miltsakaki et al., 2008; Webber et al., 2019). The hierarchy proceeds from broad classes at Level-1 (e.g., *Temporal*, *Contingency*, *Comparison*, *Expansion*) to finer types at Level-2 (e.g., *Temporal.Asynchronous*) and to specific subtypes at Level-3 (e.g., *Temporal.Asynchronous.Precedence*, as in “She finished her homework *before* dinner”). Math-DB adapts this paradigm to the mathematical domain, redefining top-level categories to align with cognitive situation schemas: Change, Combine, Compare, and Equalize. These schemas are drawn from a long tradition in mathematics education research showing that elementary arithmetic word problems cluster into a small number of recurring semantic structures, and that the structure of a problem (rather than its surface wording or arithmetic difficulty) is a primary determinant of how children represent and solve it (Carpenter and Moser, 1979; Riley et al., 1984; Fuson, 2012; Greeno, 1978; Nesher, 2020; Vergnaud, 2020). We adopt these schemas as Math-DB’s Level-1 categories because they provide a cognitively grounded inventory that maps directly onto the arithmetic operations needed to solve a problem.

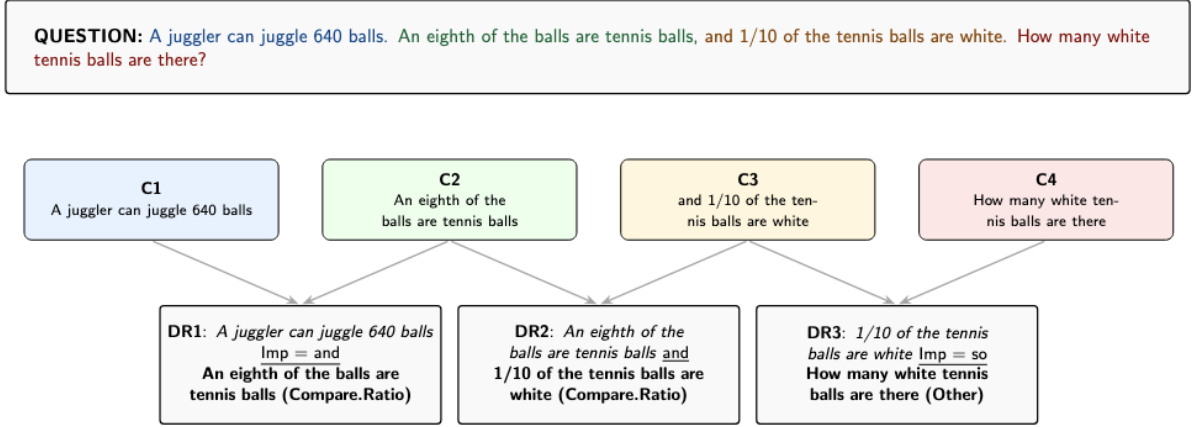


Figure 1: A schematic annotation example from Math-DB. **C** denotes a clause (component) extracted from the problem text; **DR** denotes a discourse relation between two adjacent clauses, labeled with a Math-DB Level-1/Level-2 sense; **Imp** indicates that the relation is implicit (no overt discourse connective is present in the text).

## 2.2 Math Word Problem Solving

Solving math word problems has been a long-standing challenge in AI, characterized by distinct paradigm shifts over the years. Early studies were predominantly rule-based, relying on manually crafted logic and pattern matching to map natural language directly to mathematical expressions (Bobrow, 1964; Fletcher, 1985). Subsequently, the field shifted toward machine learning-based approaches, which treated problem texts as templates to be matched to known equation forms through statistical mapping rather than hand-coded rules (Hosseini et al., 2014; Kushman et al., 2014; Roy and Roth, 2018; Shi et al., 2015).

## 2.3 LLM Reasoning and Prompting Strategies

LLMs with Chain-of-Thought (CoT) prompting offer a new way to tackle word problems, shifting from specialized pre-trained architectures (Liang et al., 2021; Shen et al., 2021) to in-context learning. CoT prompting enabled models like GPT-3/PaLM to reach state-of-the-art on GSM8K (Wei et al., 2022; Kojima et al., 2022), spurring strategies such as Least-to-Most (Zhou et al., 2022), self-consistency (Wang et al., 2022), and Tree-of-Thoughts (Yao et al., 2023).

However, LLM reasoning remains limited, particularly regarding input sensitivity. Chen et al. (2024) found that reordering premises without changing content causes significant accuracy swings. This implies LLMs lack internal mechanisms to organize information independently, relying instead on the provided narrative flow.

## 2.4 Discourse Structure for Reasoning

Recently, discourse structure has been explored as a catalyst for reasoning. Sharma et al. (2025) showed that providing general discourse context boosts GSM8K performance. Our work differs in three ways. First, while Sharma et al. (2025) use the tree-based Rhetorical Structure Theory (RST) (Mann and Thompson, 1987), Math-DB adopts the lexically-grounded PDTB paradigm (Prasad et al., 2017) for more granular analysis. Second, while Sharma et al. (2025) use general relations, Math-DB introduces a domain-specific hierarchy (e.g., Change, Combine, Compare) engineered for quantitative reasoning. Finally, unlike their synthetic annotations from a 70B model (Dubey et al., 2024), Math-DB is human-annotated and expert-verified.

In summary, literature highlights the need for structure in reliable reasoning. Our work builds on discourse analysis and LLM prompting to develop models that produce grounded, generalizable, and trustworthy reasoning.

## 3 Math-DB Framework

Math-DB provides discourse relation senses for math word problems in a two-level hierarchy (Level-1 categories, each with Level-2 subtypes). Each relation describes how Arg-2 relates to Arg-1 in a way that informs quantitative reasoning; labeling each adjacent clause pair yields a structured recipe of the reasoning steps.

The five Level-1 categories (Change, Combine, Compare, Equalize, and Other) were directly adopted from the canonical schema typology established in the mathematics education literature on

additive and subtractive word problems (Riley et al., 1984; Carpenter and Moser, 1979; Fuson, 2012). The Level-2 subtypes were derived through iterative analysis of GSM8K-style problems during a pilot annotation phase: annotators labeled a shared set of 50 problems, and recurring patterns of ambiguity and distinct semantic operations within each top-level category (e.g., additive vs. multiplicative composition under Combine) were promoted to formal subtypes. The resulting two-level hierarchy was then frozen prior to the main annotation effort.

### 3.1 Level-1 Relation Categories

We define five top-level relation categories in the Math-DB, which correspond to high-level situation schemas common in math narratives:

**Change:** Relations where Arg-2 represents an update to a quantity mentioned in Arg-1, usually through some event over time. These are the classic “change” situations (increase or decrease of some amount). For example, “Alice had 10 apples. Then she bought 5 more.”, the second sentence changes Alice’s apple count from the first sentence. Change relations involve the same quantity evolving, as opposed to combining different entities or comparing separate entities.

**Combine:** Relations where Arg-2, together with Arg-1, contributes to forming a whole. This often corresponds to part-whole addition or to the components of a multiplication scenario. Combine captures the idea of composition.

**Compare:** Relations where Arg-2 introduces a comparison between two quantities (often one from Arg-1 and another from Arg-2) via either a difference or a ratio. Unlike Combine, a comparison does not merge the quantities into a single whole, but relates them to each other (how much more/less? how many times larger/smaller?).

**Equalize:** Relations where Arg-2 is about making two quantities equal (or examining what it takes to make them equal). This often appears in problems asking “how many more... to have the same...” or scaling one quantity to match another. Equalize situations imply a goal of matching values, either by adding/removing a difference or by scaling via a ratio.

**Other:** A catch-all category for relations that do not directly contribute to the quantitative solution by introducing a new operation. We use Other in two principal cases. First, when Arg-2 provides background or contextual detail about Arg-1 that sets the scene or adds story flavor (e.g. “It was

John’s birthday.” preceding the actual math events) without introducing a new equation. Second, and equally important, when Arg-2 is the final question clause of the problem (e.g. “How many white tennis balls are there?” in Figure 1). The question clause is highly salient for problem solving, since it identifies the target quantity to be computed, but it does not itself introduce a new arithmetic operation between quantities. By labeling it as Other, we distinguish it from operation-bearing relations (Change, Combine, Compare, Equalize) while still flagging it as a structurally distinct unit in the discourse.

### 3.2 Level-2 Relation Subtypes

Each top-level category (except Other, which has no subtypes) is further refined into Level-2 senses capturing the core semantic operation or relation type. We list them below:

**Change.Increase:** Arg-2 increases the quantity from Arg-1. This corresponds to events like gaining, adding, receiving, or any scenario where the net effect is an increase in the amount. For example: “She had 5 apples. She bought 3 more.” would be Change.Increase (the 3 more increases the count).

**Change.Decrease:** Arg-2 decreases the quantity from Arg-1. Events like losing, spending, giving away, or any removal cause a decrease. E.g.: “He had 12 tickets. He gave 4 to his friend.” is Change.Decrease (4 is subtracted from 12).

Both increase and decrease are about the same item/count evolving. Often, these imply an equation  $\text{Arg-1}_{\text{quantity}} \pm \text{change} = \text{new}_{\text{quantity}}$ . If explicit, Arg-2 may contain words like “more” or “left” or “remaining” that hint at the change.

**Combine.PartWhole:** Arg-1 and Arg-2 describe parts and/or a whole in an additive part-whole structure. Typically, either:

Arg-1 and Arg-2 are two parts contributing to a total, or Arg-1 is a whole and Arg-2 gives a missing part.

For instance: “There are 7 red marbles. There are 5 blue marbles.”, Arg-2 (blue marbles) is another part to combine with Arg-1 (red marbles) into a total (implicitly 12 marbles). If later a question asks for total marbles, the relation between the two sentences is Combine.PartWhole. This subtype covers classic “aggregation” statements (“in total”, “altogether” often signal this).

**Combine.Product:** Arg-1 and Arg-2 together set up a multiplicative situation, like two factors or a structured array. A common scenario is one

sentence gives a number of groups/items per group, and the other gives another dimension (number of groups), so the total is a product. Example: “Each box has 6 eggs. There are 4 boxes.”, Arg-2 relates to Arg-1 as `Combine.Product`.

**Compare.Difference:** Arg-2 introduces an additive comparison, i.e. a difference. Typically phrased as “how many more/less” or stating one is  $N$  more than the other. For example: “Alice has 8 candies. Bob has 5 candies.” If followed by “How many more does Alice have than Bob?”, the relation between the two statements is `Compare.Difference` (Alice’s count vs Bob’s count).

**Compare.Ratio:** Arg-2 introduces a multiplicative comparison or a rate. Keywords include “times as many”, “twice/half”, or per-unit rates (“per”, “each”). For instance: “Tom has 3 marbles. Jerry has twice as many marbles as Tom.”, the second sentence is `Compare.Ratio` (Jerry’s count is a multiple of Tom’s).

**Equalize.MatchDifference:** Arg-2 frames a difference between two quantities as the adjustment needed to reach equality. Compare this to `Compare.Difference`: “Alice has 8 candies. Bob has 5. How many more does Alice have?” (the answer is a measurement of the gap) versus “Alice has 8 candies. Bob has 5. How many more does Bob need to have the same as Alice?” (the answer is the same number, 3, but interpreted as the action required to equalize). The arithmetic is identical; the discourse intent is not.

**Equalize.MatchRatio:** Arg-2 is about making quantities equal by scaling one of them, a multiplicative adjustment. E.g.: “John has 4 apples, which is twice as many as Jim. How many apples does Jim need to have the same as John?” Here, one could interpret it as Jim’s amount must be scaled up to John’s. <sup>1</sup>

---

<sup>1</sup>Both `Compare.Difference` and `Equalize.MatchDifference` involve subtraction between two quantities, and similarly both `Compare.Ratio` and `Equalize.MatchRatio` involve a multiplicative relationship. The distinction is not arithmetic but pragmatic: `Compare` relations *report* the gap or ratio between two independent quantities that remain distinct (“How many more does Alice have than Bob?”), whereas `Equalize` relations *frame the gap or ratio as an adjustment* required to reach a target state of equality (“How many more does Bob need to have the same as Alice?”). In other words, `Compare` stops at measurement, while `Equalize` prescribes an action toward equality. This pragmatic distinction matters for LLM guidance because the implied solution step differs: a `Compare` prompts a difference computation as the final answer, whereas an `Equalize` prompts the same computation but reinterprets it as “the amount to add to the smaller quantity.”

## 4 Annotation of GSM-Symbolic with Math-DB

We applied the Math-DB framework to the GSM-Symbolic benchmark, which comprises 12,500 problem instances derived from 100 GSM8K source templates. Of these, 11,414 instances passed our semi-automatic pipeline and received Math-DB sense labels, yielding 47,815 discourse relations; the remaining 1,086 instances are documented in the released diagnostic logs (see §4.4). The full annotated corpus, with all 47,815 discourse relations and accompanying metadata, is released publicly.<sup>2</sup> This section describes the dataset, our annotation procedure, corpus statistics, and pipeline-yield analysis.

### 4.1 Data

GSM-Symbolic was introduced by Mirzadeh et al. (2024) as an improved benchmark for mathematical reasoning. It was implemented by creating symbolic templates from GSM8K problems and then generating multiple instantiations by varying numbers, names, and other surface details. The goal was to test LLMs on the same underlying problems but with different wordings, to see if they truly understood the logic or just memorized specific wordings. The dataset is divided into:

**Symbolic (Main):** 5,000 problems that mirror original GSM8K questions in structure (50 instantiations per template), of which 4,621 were successfully annotated by our pipeline.

**GSM-P1 (Symbolic-Plus-One):** 5,000 problems where an additional clause has been inserted into each problem to increase complexity by one extra sentence, of which 4,535 were successfully annotated.

**GSM-P2 (Symbolic-Plus-Two):** 2,500 problems where two clauses were inserted, making the problems longer to confuse a solver, of which 2,258 were successfully annotated.

### 4.2 Annotation Procedure

Before describing the pipeline, we note an important property of GSM-Symbolic that shaped our annotation effort. Because instances within a given template differ only in surface details (numerical values, proper names, units) and not in semantic structure, all instances of a single template share the same underlying discourse structure. We

---

<sup>2</sup><https://github.com/erolcan-er/Math-DB>

therefore performed template-level (id-based) annotation for the main GSM-Symbolic subset: one canonical sense sequence was annotated per template and propagated to instances of that template for which our Stage-1 pipeline could re-apply the canonical clause boundaries to the variant surface form. Instances for which canonical alignment failed (i.e., the detected number of discourse relations did not match the template) were excluded and are documented in the diagnostic logs (§4.4). This reduced the annotation effort on the main subset from approximately 21k relations to a much smaller set of unique templates.

However, for GSM-P1 and GSM-P2, we adopted an instance-level annotation strategy. Although surface variation within a template is still semantic-preserving, the inserted clauses in P1 and P2 are themselves generated from a template, so different instances of the same id can include structurally different inserted clauses (e.g., a distractor in one instance versus a real additional reasoning step in another). We therefore annotated each P1/P2 instance independently to capture this variation; in practice, different instances of the same P1/P2 id can differ in both the number and the sense of their discourse relations. The corpus statistics in Table 1 are reported over all 47,815 final relations across all instances (main, P1, and P2) rather than over unique templates.

We used a two-stage annotation strategy: automatic pre-processing followed by manual sense labeling.

#### 4.2.1 Stage 1: Connective Detection and Argument Segmentation

We fine-tuned BERT large (Devlin et al., 2019) using a sub dataset of GSM8K labeled by BIO tags (following DISRPT shared task formulations (Zeldes et al., 2021; Braud et al., 2023)) to identify explicit discourse connectives in the text (words like “then”, “so”, “after”, “because”, etc.) and to segment sentences into two arguments (Arg-1 and Arg-2) around each connective. If two arguments are detected without explicit discourse connectives, we accept these arguments as an implicit discourse relation<sup>3</sup>.

This step provided discourse relations consist of Arg-1 and Arg-2 spans together with optional discourse connective spans for explicit relations.

<sup>3</sup>See Appendix A for detailed model specifications and training hyperparameters.

#### 4.2.2 Stage 2: Sense Labeling by Annotators

For the Stage 2, we recruited six undergraduate students as annotators from the mathematics department (fluent in English and with strong math aptitude). Each discourse relations were assigned to three annotators independently for sense labeling. Annotators were provided with the Math-DB annotation manual (as described in Appendix B, with examples and decision criteria). For each discourse relation, they choose the appropriate Level-1 and Level-2 sense that best describes how Arg-2 relates to Arg-1 in the context of solving the problem. We allowed annotators to see the full problem text (all sentences) so they had context, but instructed them to label relations in a forward sequential manner, treating each adjacent link in turn.

To ensure quality, we conducted an initial training phase where all annotators labeled a common set of 50 problems, then discussed disagreements with an adjudicator (the expert author). This helped clarify guidelines (e.g., distinguishing Combine vs. Change when an addition happens over time: if a time/event is involved, it is Change, but if it is just combining two existing quantities, it is Combine). After refining the instructions, the annotators proceeded with the full dataset.

Each relation was assigned to three of the six annotators (sampled to balance load), so that every relation ultimately received three independent labels. We aggregated them by majority vote: if at least two annotators agreed on the same sense, that was assigned as the final label.

For cases where all three annotators gave different labels or one label was not in majority (this happened for 1,342 relations, 2.8% of the total), we invoked the expert adjudication. The expert (a computational linguist with a strong math background, also an author) reviewed those cases and chose the final sense label among the proposed ones. After adjudication, every relation in the dataset had a final agreed-upon label.

### 4.3 Corpus Statistics and Analysis

Our Math-DB annotated corpus comprises 47,815 discourse relations across 11,414 successfully-aligned instances (out of 12,500 source problems; the remaining 1,086 are analyzed in §4.4).

**Explicit vs Implicit:** Of the 47,815 relations, 15,011 ( $\approx 31.4\%$ ) are explicit discourse relations. The remaining 32,804 ( $\approx 68.6\%$ ) are implicit discourse relations, meaning no connective is overtly

present and the relation is assumed between each adjacent clause. This implicit majority is expected, since many math problem sentences follow one another without “so/then/because” even though a logical link exists. It underscores the need for reasoning beyond surface cues.

Relation Sense	Count	%
<b>Change</b>	<b>14,891</b>	<b>31.1%</b>
Change.Increase	8,412	17.6%
Change.Decrease	6,479	13.5%
<b>Combine</b>	<b>7,484</b>	<b>15.7%</b>
Combine.PartWhole	5,401	11.3%
Combine.Product	2,083	4.4%
<b>Compare</b>	<b>1,498</b>	<b>3.1%</b>
Compare.Ratio	1,020	2.1%
Compare.Difference	478	1.0%
<b>Equalize</b>	<b>504</b>	<b>1.1%</b>
Equalize.MatchDifference	430	0.9%
Equalize.MatchRatio	74	0.2%
<b>Other</b>	<b>23,438</b>	<b>49.0%</b>
<b>Total</b>	<b>47,815</b>	<b>100.0%</b>

Table 1: Distribution of Math-DB discourse relations across Level-1 categories and Level-2 subtypes, computed over 47,815 relations from 11,414 successfully-annotated instances.

**Distribution of Level-1 senses:** Table 1 reports the full sense distribution. The single largest category is *Other* at 49.0%, reflecting that nearly half of all adjacent clause pairs are either question clauses or narrative-context links that do not introduce a new arithmetic operation. Among the operation-bearing categories, *Change* is the most frequent (31.1%), followed by *Combine* (15.7%); together they account for roughly 47% of all relations, consistent with the fact that most problems involve accumulating or modifying a running total (Change) and aggregating pieces (Combine). *Compare* (3.1%) and *Equalize* (1.1%) are substantially less frequent, which aligns with the GSM-style problem inventory: comparison and equalization scenarios appear but are not the dominant problem types.

**Level-2 senses:** Within Change, increases (8,412) outnumber decreases (6,479), reflecting that problem narratives more often involve gaining (earning, buying, receiving) before possibly losing or spending later. Within Combine, PartWhole (5,401) is roughly 2.6 $\times$  more frequent than Product (2,083), indicating that pure additive composition is more common than explicit multiplicative composition. For Compare, Ratio (1,020) outnumbers

Difference (478) by roughly 2 $\times$ , which is consistent with the prevalence of unit-rate and “times as many” constructions in GSM-style problems. For Equalize, MatchDifference (430) is far more common than MatchRatio (74), aligning with typical curricula in which “how many more to have same” scenarios appear more often than scale-equalize scenarios.

We also tracked the occurrence of connectives. Some explicit connectives we saw frequently: temporal connectives like “then”, “after that” signaling Change relations, additive connectives like “in total” or “altogether” signaling Combine.PartWhole, contrastive phrases like “more than” or “less than” indicating Compare.Difference, phrases like “times as many” for Compare.Ratio, and equality phrases like “the same as” for Equalize. We note that our connective inventory is open-ended: the span-based Stage-1 detector occasionally identifies multi-token linking phrases (e.g., “On top of that,” “each having”) as connectives. We retain these as found rather than collapsing them to a closed PDTB-style lexicon, since they often carry semantic content relevant to the relation.

Inter-annotator agreement before adjudication was substantial: Fleiss’  $\kappa = 0.92$  on Level-1 senses and  $\kappa = 0.84$  on Level-2 senses (full IAA details and confusion analysis in Appendix D).

#### 4.4 Pipeline Yield and Failure Modes

Of the 12,500 GSM-Symbolic instances we set out to annotate, our semi-automatic pipeline successfully produced Math-DB labels for 11,414 instances (91.3%). The remaining 1,086 instances were rejected by the alignment stage with the diagnostic reason `dr_count_mismatch`: the number of discourse relations detected in the variant did not match the canonical template’s expected count, indicating that the variant’s surface form deviated from the template structure in ways our segmenter could not reconcile. Table 2 reports the per-subset yield.

Subset	Total	Annotated	Yield
Symbolic (Main)	5,000	4,621	92.4%
GSM-P1	5,000	4,535	90.7%
GSM-P2	2,500	2,258	90.3%
<b>Total</b>	<b>12,500</b>	<b>11,414</b>	<b>91.3%</b>

Table 2: Per-subset pipeline yield for Math-DB annotation. Rejected instances are documented in the released diagnostic logs.

We release the full diagnostic logs (instance identifiers, detected vs. expected relation counts, and failure reasons) alongside the annotated corpus. These logs serve a dual purpose: they make pipeline coverage transparent to downstream users, and they provide a targeted benchmark for future work on robust automatic discourse segmentation of mathematical text.

## 5 Using Discourse Annotations to Enhance CoT Prompting

In this section, we leverage our annotations to improve LLM reasoning. Our approach operates in the context of CoT prompting. The key idea is to supply the LLM with Math-DB guided hints about what each step of the reasoning should do, in addition to the problem text itself.

### 5.1 Discourse-Augmented Prompt Design

All experiments in this section use GPT-4o-mini as the underlying LLM, chosen to match the model evaluated in the original GSM-Symbolic study (Mirzadeh et al., 2024) for direct comparability. We discuss the implications of this single-model setup, and the expected behavior across other LLMs and prompting strategies, in the Limitations section.

A straightforward way to include discourse information is to present it as an annotated version of the problem. We designed a prompting format that uses a structured format with discourse sense information for each targeted problem.

For each targeted problem, we present a brief instruction for Math-DB annotation scheme. Then, we append 8-shot demonstrations with Math-DB discourse relations and step-by-step answer explanations. In the inference phase, discourse relations found in the targeted problem are provided. Finally, we ask the model to solve the given problem using all provided information. Simplified prompt template can be seen in Figure 2<sup>4</sup>:

For our baseline, we follow the standard evaluation protocol defined in the Mirzadeh et al. (2024). The baseline uses an 8-shot CoT prompt using the original GSM8K demonstrations. The prompt template consists of a general system instruction followed by eight fixed question-answer pairs that provide only the problem text and the step-by-step reasoning chain, without any discourse-level annotations. All experiments were conducted using greedy decoding to ensure consistency. The full

<sup>4</sup>The full prompt template is provided in Appendix C.

prompt template for the baseline approach is provided in Appendix C.

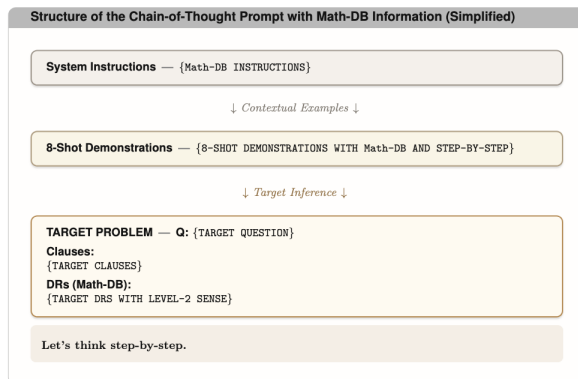


Figure 2: Simplified prompt template

In our evaluation, the discourse relations provided to the model during the inference phase are derived from the gold-standard annotations in Math-DB. We intentionally adopt this ‘‘Oracle’’ experimental setup to isolate the impact of the discourse framework on mathematical reasoning from the potential noise and error propagation of an automated parser. By using ground-truth relations, we establish the upper-bound performance gain achievable when the underlying semantic structure of a problem is explicitly represented for the LLM.

### 5.2 Results

Our discourse-augmented prompting consistently outperforms the standard CoT prompting on all subsets of GSM-Symbolic. We provide a Figure 3 of accuracies below.

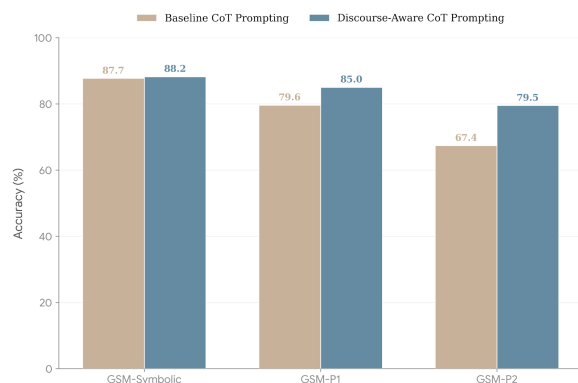


Figure 3: Comparison of accuracy between Baseline CoT Prompting and Discourse-Aware CoT Prompting across three GSM-based datasets.

On the GSM-Symbolic dataset, GPT-4o-mini with standard CoT achieved an accuracy of 87.7%,

whereas with Math-DB prompting it achieved 88.2%.

On GSM-P1, the baseline performance dropped to 79.6% (highlighting the brittleness when an extra irrelevant clause is present, similar to observations from Mirzadeh et al. (2024)). In contrast, our method achieved 85% on P1, which is substantially higher, and notably close to the Symbolic dataset performance. The relative drop from Symbolic  $\rightarrow$  P1 was much smaller with our approach compared to the baseline CoT prompting strategy.

On GSM-P2, the trend is similar. Baseline CoT struggled even more with two distractors, achieving 67.4%, whereas our discourse-guided approach obtained 79.5%. While GSM-P2 is hardest for both, the gap between our method and baseline widened further here, showing that the more complex the input, the more the structured guidance helps.

### 5.3 Error Analysis

We performed a qualitative analysis of cases where the baseline CoT failed but the discourse-augmented CoT succeeded, and vice versa.<sup>5</sup>

Baseline CoT often failed on P1/P2 datasets with following reasons:

- getting distracted by a fake number or event (e.g., trying to use a number from an irrelevant clause in the calculation).
- losing track of the problem structure (like forgetting a previous value or mixing up which numbers to add vs subtract).

With Math-DB prompts, such errors were moderately reduced. We saw the model explicitly ignore irrelevant sentences labeled “Other” in its solution, which is precisely what we want it to do. We also observed that sense labels disambiguated keyword-driven errors. For instance, the keyword “more” no longer reliably triggered addition in clauses labeled Compare.Difference. This indicates that the model correctly used the relation information both to filter out distractions and to select the intended arithmetic operation.

<sup>5</sup>While the error analysis provided here is qualitative, it is intended to illustrate one potential application of the Math-DB framework. Specifically, its utility in enhancing prompt-based reasoning rather than to serve as an exhaustive evaluation of LLM performance. The primary contribution of this work remains the definition, annotation, and validation of the Math-DB schema itself.

### 5.4 Discussion

The gains stem from two effects. First, “Other” labels prune the search space by explicitly marking clauses to exclude from the math, preventing distractor errors. Second, sense labels disambiguate operation-triggering keywords (e.g. “more” under Compare.Difference no longer defaults to addition), reducing logical errors. LLMs benefit from explicit guidance that plays to their strength in pattern-following.

## 6 Conclusion and Future Work

We introduced Math-DB, a domain-specific discourse framework and annotated corpus for mathematical word problems, and showed that Math-DB annotations enable GPT-4o-mini to maintain high accuracy on perturbed GSM-Symbolic problems, where standard CoT degrades. Future work includes:

**Generality to Other Datasets:** We focused on GSM-Symbolic. It would be interesting to apply discourse guided prompting to other math problem sets (like the original GSM8K, or more complex ones like MAWPS or Math23K, etc.). The relations are quite general for basic arithmetic word problems, but extending the taxonomy (e.g., for multi-step algebra or physics problems) could be needed.

**Broader Reasoning Domains:** The idea of domain-specific discourse frameworks could be applied beyond math. Other reasoning tasks (like scientific question answering, logical deduction puzzles, or multi-hop reading comprehension) might benefit from a similar approach.

**End-to-End Pipeline:** Our experiments use gold-standard Math-DB annotations as a proof of concept, establishing the upper-bound benefit of structured guidance. A natural next step is to build a fully autonomous parser-to-LLM pipeline in which discourse relations are predicted automatically at inference time and fed directly into the reasoning model.

In conclusion, the Math-DB demonstrates a synergy between discourse theory and LLMs.

### Limitations

While Math-DB demonstrates clear benefits for mathematical reasoning, several limitations for this study should be considered:

- **Dataset Specificity:** This study focused primarily on the GSM-Symbolic dataset. While

the discourse relations are generalizable to basic arithmetic word problems, their applicability to more advanced domains, such as multi-step algebra, calculus, or physics may require further refinement and extension of the existing taxonomy.

- **Annotation Intensity and Scaling:** The current dataset was developed using a semi-automatic pipeline involving manual sense labeling by six trained annotators with mathematics backgrounds. While this ensured high-quality data, the reliance on human expertise for labeling represents a potential bottleneck for scaling the framework to significantly larger or more diverse datasets. A promising direction is to leverage the existing 47k annotated relations as training data for a supervised sense classifier, which could then propagate labels to new problems via bootstrapping or active learning, with human annotators verifying only the low-confidence cases. We leave the development and evaluation of such a scalable annotation pipeline to future work.
- **Pipeline Dependence and Segmentation Errors:** The annotation pipeline relies on an automatic preprocessing stage (connective detection and argument segmentation with a fine-tuned BERT model), followed by manual sense labeling. Any systematic segmentation errors (e.g., incorrect clause boundaries or misidentified connectives) can affect the validity of the human annotation process (sense labeling) negatively.
- **Oracle Evaluation and Pipeline Dependencies:** Our experiments utilize ground-truth discourse labels to evaluate the framework’s reasoning potential, characterizing our results as an upper-bound analysis rather than a system-wide evaluation of a fully automated pipeline. In a practical deployment, any systematic errors from the initial segmentation or labeling stages (as discussed in Appendix A) would propagate to the reasoning model, potentially degrading performance. Future work is required to develop high-accuracy automated discourse parsers that can bridge this gap for end-to-end applications.
- **Focus on Adjacent Relations:** Our current framework primarily annotates adjacent

clause pairs, as long-distance links are relatively rare in the short problems found in GSM-Symbolic, GSM-P1 and GSM-P2. However, more complex, multi-paragraph math problems may require capturing longer-distance discourse dependencies that the current sequential approach might overlook.

- **Language Scope:** The annotated corpus and experiments were conducted solely in English. Future work is needed to determine the framework’s effectiveness across different languages to varied linguistic structures in a multilingual context.
- **Pipeline Coverage:** Our semi-automatic pipeline successfully annotated 11,414 of 12,500 GSM-Symbolic instances (91.3%); the remaining 1,086 instances failed at the alignment stage with detected discourse-relation counts inconsistent with their canonical templates. While we release full diagnostic logs for these instances, improving the robustness of the Stage-1 segmenter to reduce this failure rate is left to future work.
- **Single-Model Evaluation:** Our experiments use GPT-4o-mini as the underlying LLM, chosen because it was the model evaluated in the original GSM-Symbolic study (Mirzadeh et al., 2024), enabling direct comparison. The interaction between Math-DB guidance and (i) larger or more capable reasoning models (e.g., GPT-4, Claude, or specialized reasoning variants), and (ii) alternative structured prompting strategies (e.g., Least-to-Most, Tree-of-Thoughts, self-consistency) remains an open question. We hypothesize that discourse guidance provides complementary signal to these strategies rather than a substitute, but empirical verification is left to future work.

## References

- Daniel G Bobrow. 1964. A question-answering system for high school algebra word problems. In *Proceedings of the October 27-29, 1964, fall joint computer conference, part I*, pages 591–614.
- Chloé Braud, Yang Janet Liu, Eleni Metheniti, Philippe Muller, Laura Rivière, Attapol Rutherford, and Amir Zeldes. 2023. The disrpt 2023 shared task on elementary discourse unit segmentation, connective detection, and relation classification. In *Proceedings*

- of the 3rd Shared Task on Discourse Relation Parsing and Treebanking (DISRPT 2023), pages 1–21.
- Thomas P Carpenter and James M Moser. 1979. An investigation of the learning of addition and subtraction.
- Xinyun Chen, Ryan A Chi, Xuezhi Wang, and Denny Zhou. 2024. Premise order matters in reasoning with large language models. *arXiv preprint arXiv:2402.08939*.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, and 1 others. 2021. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*.
- Gabriella Daroczy, Magdalena Wolska, Walt Detmar Meurers, and Hans-Christoph Nuerk. 2015. Word problems: A review of linguistic and numerical factors contributing to their difficulty. *Frontiers in psychology*, 6:348.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 conference of the North American chapter of the association for computational linguistics: human language technologies, volume 1 (long and short papers)*, pages 4171–4186.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, and 1 others. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Charles R Fletcher. 1985. Understanding and solving arithmetic word problems: A computer simulation. *Behavior Research Methods, Instruments, & Computers*, 17(5):565–571.
- Karen C Fuson. 2012. *Children’s counting and concepts of number*. Springer Science & Business Media.
- James G Greeno. 1978. Natures of problem-solving abilities. In *Handbook of learning and cognitive processes: Vol. 5. Human information processing*, pages 240–269. Erlbaum Hillsdale, NJ.
- Mohammad Javad Hosseini, Hannaneh Hajishirzi, Oren Etzioni, and Nate Kushman. 2014. Learning to solve arithmetic word problems with verb categorization. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pages 523–533.
- Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. 2022. Large language models are zero-shot reasoners. *Advances in neural information processing systems*, 35:22199–22213.
- Nate Kushman, Yoav Artzi, Luke Zettlemoyer, and Regina Barzilay. 2014. Learning to automatically solve algebra word problems. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 271–281.
- Zhenwen Liang, Jipeng Zhang, Jie Shao, and Xiangliang Zhang. 2021. Mwp-bert: A strong baseline for math word problems.
- William C Mann and Sandra A Thompson. 1987. Rhetorical structure theory: Description and construction of text structures. In *Natural language generation: New results in artificial intelligence, psychology and linguistics*, pages 85–95. Springer.
- Eleni Miltsakaki, Livio Robaldo, Alan Lee, and Aravind Joshi. 2008. Sense annotation in the penn discourse treebank. In *International Conference on Intelligent Text Processing and Computational Linguistics*, pages 275–286. Springer.
- Iman Mirzadeh, Keivan Alizadeh, Hooman Shahrokhi, Oncel Tuzel, Samy Bengio, and Mehrdad Farajtabar. 2024. Gsm-symbolic: Understanding the limitations of mathematical reasoning in large language models. *arXiv preprint arXiv:2410.05229*.
- Pearla Neshet. 2020. Levels of description in the analysis of addition and subtraction word problems. In *Addition and subtraction*, pages 25–38. Routledge.
- Rashmi Prasad, Bonnie Webber, and Aravind Joshi. 2017. The penn discourse treebank: An annotated corpus of discourse relations. In *Handbook of linguistic annotation*, pages 1197–1217. Springer.
- Peng Qi, Yuhao Zhang, Yuhui Zhang, Jason Bolton, and Christopher D Manning. 2020. Stanza: A python natural language processing toolkit for many human languages. *arXiv preprint arXiv:2003.07082*.
- Mary S Riley and 1 others. 1984. Development of children’s problem-solving ability in arithmetic.
- Subhro Roy and Dan Roth. 2018. Mapping to declarative knowledge for word problem solving. *Transactions of the Association for Computational Linguistics*, 6:159–172.
- Krish Sharma, Niyar R Barman, Akshay Chaturvedi, and Nicholas Asher. 2025. Dimsum: Discourse in mathematical reasoning as a supervision module. *arXiv preprint arXiv:2503.04685*.
- Jia Tracy Shen, Michiharu Yamashita, Ethan Prihar, Neil Heffernan, Xintao Wu, Ben Graff, and Dongwon Lee. 2021. Mathbert: A pre-trained language model for general nlp tasks in mathematics education. *arXiv preprint arXiv:2106.07340*.
- Freda Shi, Xinyun Chen, Kanishka Misra, Nathan Scales, David Dohan, Ed H Chi, Nathanael Schärli, and Denny Zhou. 2023. Large language models can be easily distracted by irrelevant context. In *International Conference on Machine Learning*, pages 31210–31227. PMLR.

- Shuming Shi, Yuehui Wang, Chin-Yew Lin, Xiaojiang Liu, and Yong Rui. 2015. Automatically solving number word problems by semantic parsing and reasoning. In *Proceedings of the 2015 conference on empirical methods in natural language processing*, pages 1132–1142.
- G rard Vergnaud. 2020. A classification of cognitive tasks and operations of thought involved in addition and subtraction problems. In *Addition and subtraction*, pages 39–59. Routledge.
- Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. 2022. Self-consistency improves chain of thought reasoning in language models. *arXiv preprint arXiv:2203.11171*.
- Bonnie Webber, Rashmi Prasad, Alan Lee, and Aravind Joshi. 2019. The penn discourse treebank 3.0 annotation manual. *Philadelphia, University of Pennsylvania*, 35:108.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, and 1 others. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837.
- Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Tom Griffiths, Yuan Cao, and Karthik Narasimhan. 2023. Tree of thoughts: Deliberate problem solving with large language models. *Advances in neural information processing systems*, 36:11809–11822.
- Amir Zeldes, Yang Janet Liu, Mikel Iruskieta, Philippe Muller, Chlo  Braud, and Sonia Badene. 2021. The disrpt 2021 shared task on elementary discourse unit segmentation, connective detection, and relation classification. In *Proceedings of the 2nd Shared Task on Discourse Relation Parsing and Treebanking (DISRPT 2021)*, pages 1–12.
- Denny Zhou, Nathanael Sch rli, Le Hou, Jason Wei, Nathan Scales, Xuezhi Wang, Dale Schuurmans, Claire Cui, Olivier Bousquet, Quoc Le, and 1 others. 2022. Least-to-most prompting enables complex reasoning in large language models. *arXiv preprint arXiv:2205.10625*.

## A Appendix: Data Preparation and Parameter Tuning for the BERT Fine-Tuning in Discourse Connective Detection and Argument Span Labeling Tasks

This appendix describes how we fine-tune a BERT-based sequence labeling model for (i) discourse connective detection and (ii) argument span labeling (Arg-1/Arg-2) in mathematical word problems (MWP). We formulate both subtasks as a single token-level named entity recognition (NER) problem using a BIO tagging scheme. Tokenization is performed with Stanza (Qi et al., 2020), and each token is assigned one label from the set

$$\mathcal{Y} = \{0, \text{B-Conn}, \text{I-Conn}, \text{B-Arg1}, \text{I-Arg1}, \text{B-Arg2}, \text{I-Arg2}\}$$

Tokens outside any discourse span are labeled 0. We use the convention that B- marks the beginning of a span and I- marks tokens inside the span.

### A.1 Task formulation as token-level NER

Given an input text (a window containing a discourse relation instance), we produce a token sequence  $(x_1, \dots, x_n)$  and a corresponding label sequence  $(y_1, \dots, y_n)$  with  $y_i \in \mathcal{Y}$ . The goal is to predict the discourse connective span and the two discourse arguments (Arg-1 and Arg-2) as contiguous token spans. This converts connective detection and argument span labeling into standard BIO sequence labeling.

**Relation-instance encoding:** A single MWP typically contains multiple adjacent clause relations (e.g.,  $\text{Clause}_1 \rightarrow \text{Clause}_2$ ,  $\text{Clause}_2 \rightarrow \text{Clause}_3$ , etc.). To avoid label conflicts in a single long sequence, we construct the dataset using *relation instances*. For each adjacent clause pair (Arg-1  $\rightarrow$  Arg-2), we create one training example consisting of a local text window that contains Arg-1, an optional connective (if explicit), and Arg-2. Each relation therefore yields one NER instance with a single Arg-1 span, a single Arg-2 span, and optionally one connective span. This representation is both simpler and empirically more stable than whole-problem multi-span tagging.

**BIO labeling rules:** For each gold span type  $S \in \{\text{Conn}, \text{Arg-1}, \text{Arg-2}\}$ , we label the first token overlapping the span as B- $S$  and any subsequent overlapping tokens as I- $S$ . Tokens that do not overlap any gold span are labeled 0.

### A.2 Illustrative example

Table 3 shows an example relation instance and its BIO labels in a CoNLL-style format. Arg-1 is the first sentence, *Then* is an explicit connective, and Arg-2 is the second sentence.

Token	Label
Lina	B-Arg1
has	I-Arg1
8	I-Arg1
marbles	I-Arg1
.	I-Arg1
Then	B-Conn
she	B-Arg2
gives	I-Arg2
3	I-Arg2
marbles	I-Arg2
to	I-Arg2
her	I-Arg2
friend	I-Arg2
.	I-Arg2

Table 3: Example of tokenization and BIO labels for a discourse relation instance.

For *implicit* relations, we do not fine-tune any model since we assume all adjacent clauses build an implicit discourse relation if they do not build an explicit discourse relation.

### A.3 Model architecture

We fine-tune a pretrained BERT encoder with a token classification head: (i) a Transformer encoder (e.g., bert-base-uncased or bert-base-cased), and (ii) a linear layer mapping contextual representations  $h_i$  to label logits over  $\mathcal{Y}$ . We optimize token-level cross-entropy loss over non-ignored positions.

### A.4 Training objective and evaluation

**Objective:** Given token sequence  $(x_1, \dots, x_n)$ , the model predicts label distribution  $p(y_i | x)$  for each token. We minimize the negative log-likelihood (cross-entropy) over labeled positions:

$$\mathcal{L} = - \sum_{i=1}^n \mathbf{1}[y_i \neq \text{IGNORE}] \log p(y_i | x).$$

**Metrics:** Table 4 reports span-level (entity-level)  $F_1$  scores for each predicted span type. The model achieves very strong performance on connective span detection ( $F_1=0.96$ ), reflecting that explicit connectives are often short and lexically distinctive. Arg-2 spans are also identified reliably ( $F_1=0.91$ ). Arg-1 span labeling is comparatively harder ( $F_1=0.86$ ), which we attribute to more variable Arg-1 boundaries and frequent adjacency to preceding context, making precise start/end decisions less consistent than for connectives or Arg-2.

Task	Precision (%)	Recall (%)	$F_1$ (%)
Connective span detection	97.11	95.50	96.30
Arg-1 span labeling	87.32	84.58	85.93
Arg-2 span labeling	92	89.78	90.88

Table 4: Span-level (entity-level) performance of the BERT token-classification model on the test set.

**Discussion:** Table 4 shows strong span extraction performance across all subtasks. Connective span detection is the easiest ( $F_1=0.96$ ), likely because explicit connectives are typically short and lexically distinctive. Arg2 spans are also identified reliably ( $F_1=0.91$ ). Arg1 span labeling is comparatively harder ( $F_1=0.86$ ), which we attribute to greater boundary variability (e.g., optional lead-in context and punctuation) and stronger dependency on discourse segmentation decisions.

### A.5 Hyperparameter tuning

We treat the task as standard NER fine-tuning and tune hyperparameters on the dev set, selecting the checkpoint with the best dev micro span-level  $F_1$ . We recommend reporting the full search space and the chosen configuration.

#### Core hyperparameters (tuned):

- Learning rate:  $\{1 \times 10^{-5}, 2 \times 10^{-5}, 3 \times 10^{-5}, 5 \times 10^{-5}\}$
- Batch size (effective):  $\{16, 32\}$
- Epochs:  $\{3, 5, 10\}$  with early stopping on dev  $F_1$
- Max sequence length:  $\{128, 256, 384\}$  (relation windows typically fit within 256)
- Warmup ratio:  $\{0.0, 0.06, 0.1\}$
- Dropout: BERT default (0.1)

We use AdamW with weight decay 0.01 and a linear learning rate schedule with decay.

**Seeds and reporting:** We run each configuration with multiple random seeds (e.g., 3) and report  $\text{mean} \pm \text{std}$  on test  $F_1$  for robustness.

## B Appendix: Math-DB Annotation Manual

### B.1 Purpose and Scope

Math-DB adapts the Penn Discourse TreeBank (PDTB) framework to math word problems.

- A **problem** is represented as an **sequence of clauses**  $C_1, C_2, \dots, C_n$  in text order.
- For every **adjacent clause pair**  $(\text{Arg-1} \rightarrow \text{Arg-2}) = (C_i \rightarrow C_{i+1})$ , annotate **exactly one** discourse relation (DR) sense.
- DR senses describe how **Arg-2 relates to Arg-1** in a math-reasoning way and are intended to scaffold quantitative reasoning steps.

Math-DB is designed for grade-school arithmetic word problems (e.g., GSM8K-style) and emphasizes relations that map directly to common quantitative operations.

### B.2 Units of Annotation

#### B.2.1 Clause

A **clause** is the smallest meaningful proposition used for math reasoning. Clauses typically contain: (i) a quantity (explicit number or implied variable), (ii) an entity (e.g., apples, money, students), and (iii) a predicate expressing a state/event/relation (e.g., “has”, “buys”, “costs”, “is”, “more than”).

Clauses are labeled in order:  $C_1, C_2, \dots, C_n$ .

#### B.2.2 Discourse Arguments (Arg-1, Arg-2)

For each adjacent pair:

$$\text{Arg-1} = C_i, \quad \text{Arg-2} = C_{i+1}.$$

The relation direction is always forward in text order (Arg-1 precedes Arg-2). Only **adjacent pairs** are annotated.

### B.3 Sense Inventory

Each discourse relation has the form `Level1.Level2`, where

$$\text{Level1} \in \{\text{Change, Combine, Compare, Equalize, Other}\}.$$

#### B.3.1 Change (State transition over time)

Arg-2 updates a quantity introduced in Arg-1 via an event over time.

- `Change.Increase`: the tracked quantity increases (gain, receive, buy, plant, add, arrive).
- `Change.Decrease`: the tracked quantity decreases (lose, spend, give away, remove, eat, sell, break).

**Diagnostic:** Arg-1 and Arg-2 refer to the **same tracked quantity** at different times.

#### B.3.2 Combine (Composition into a whole)

Arg-2 contributes to forming a whole, either additively or multiplicatively.

- `Combine.PartWhole`: aggregation of parts into a total (*total, altogether, combined, in all, the rest*).
- `Combine.Product`: structured multiplication ( $m \times n$ ), including arrays/area, repeated groups, combinations, and unit-rate multiplication (“\$3 each”, “5 per day”).

**Diagnostic:** Arg-1 and Arg-2 jointly define a **composed quantity** (sum or product), not a time update.

### B.3.3 Compare (Relational linking without merging)

Arg-2 relates two quantities by difference or ratio without forming a single whole.

- Compare.Difference: additive comparison (“how many more/less”, difference).
- Compare.Ratio: multiplicative comparison or rate (“times as many”, “per”, “each”, unit price/speed/density).

**Diagnostic:** Two quantities are linked for comparison, not combined into one total.

### B.3.4 Equalize (Reach equality)

Arg-2 frames a goal of making quantities match, either by additive adjustment or scaling.

- Equalize.MatchDifference: reach equality by adding/removing an amount (“how many more needed to be the same”).
- Equalize.MatchRatio: reach equality by scaling (“make it  $k$  times”, “scale until equal”).

**Diagnostic:** The intended target state is **equality**.

### B.3.5 Other (Non-quantitative or auxiliary)

Use Other when the clause pair does not introduce a new quantitative step or core math relation. This covers two main cases: (i) narrative context or descriptive details not used in the computation, and (ii) the final question clause of the problem, which identifies the target quantity but does not itself introduce an arithmetic operation between quantities. Prefer a non-Other sense when Arg-2 changes, composes, compares, or equalizes quantities in a way used for solving.

## B.4 Clause Segmentation Guidelines

### B.4.1 General rule

Split into clauses so that each clause expresses **one main proposition** relevant to the math.

### B.4.2 Strong split cues

Split at:

- sentence boundaries,
- semicolons,
- discourse markers (*then, after, later, before, when, while, however, but*) when they introduce a new event or relation,
- enumerations that introduce separate quantities.

### B.4.3 Keep together when

Do **not** split when a phrase is tightly bound as one atomic math fact (e.g., “5 bagels for \$3 each”), unless separate handling is needed.

### B.4.4 The question

The final interrogative typically becomes its own clause ( $C_n$ ).

## B.5 Relation Selection Heuristics

Use these diagnostics in order:

1. Does Arg-2 **update a previously tracked quantity over time**? → Change.Increase or Change.Decrease
2. Does Arg-2 **compose a whole** from parts or groups? → Combine.PartWhole or Combine.Product
3. Does Arg-2 **compare** quantities (difference or ratio) without merging them? → Compare.Difference or Compare.Ratio
4. Does Arg-2 express **reaching equality**? → Equalize.MatchDifference or Equalize.MatchRatio
5. Otherwise → Other

## B.6 Common Constructions and Recommended Labels

### B.6.1 “Each / per / at a rate”

- If used to compute a total (price  $\times$  quantity, speed  $\times$  time): Combine.Product.
- If emphasized as a relational comparison: Compare.Ratio.
- When uncertain in GSM-style problems, prefer Combine.Product for “each/per” totals.

### B.6.2 “X more / fewer / less than”

- Usually Compare.Difference.
- If defining a required adjustment to match a target: Equalize.MatchDifference.

### B.6.3 “Times as many / double / triple / half”

- If comparing two quantities: Compare.Ratio.
- If computing a produced total across repeated groups: Combine.Product.
- “Half ... swam away” is typically Change.Decrease (event reducing count).

### B.6.4 Multi-step updates

Successive events on the same quantity are annotated as a chain of Change.\* relations across adjacent clauses.

### B.6.5 “The rest”

Often Combine.PartWhole (whole/part inference), possibly followed by Change.Decrease if phrased as removal.

## B.7 Worked Example

**Problem:** “A juggler can juggle 640 balls. An eighth of the balls are tennis balls, and 1/10 of the tennis balls are white. How many white tennis balls are there?”

### Clauses

- $C_1$ : A juggler can juggle 640 balls.
- $C_2$ : An eighth of the balls are tennis balls.
- $C_3$ : 1/10 of the tennis balls are white.
- $C_4$ : How many white tennis balls are there?

**DRs**

- Arg-1 =  $C_1 \rightarrow$  Arg-2 =  $C_2$ : Combine.Product
- Arg-1 =  $C_2 \rightarrow$  Arg-2 =  $C_3$ : Combine.Product
- Arg-1 =  $C_3 \rightarrow$  Arg-2 =  $C_4$ : Other

## C Appendix: Prompt Templates for CoT Prompting with Math-DB and Baseline Study

### Prompt Template: Math-DB Reasoning

#### INSTRUCTIONS:

You are a mathematical reasoning engine. Solve the question step by step.

#### A) Mathematical DB Framework (Math-DB)

Some problems may come with a DB-style sequence of discourse relations (DR) between adjacent clauses. Each DR is a label of the form Level1.Level2 describing how Arg-2 relates to Arg-1. Level-1 senses:

- Change: Arg-2 updates a quantity from Arg-1 via an event.
- Combine: Arg-1 and Arg-2 compose into a whole (additively or via a structured product).
- Compare: Arg-2 relates two quantities by difference or by ratio/rate.
- Equalize: Arg-2 expresses making quantities match (equality), by difference or by scaling.
- Other: adjacency without a quantitative relation from this inventory.

Level-2 senses:

- Change.Increase / Change.Decrease
- Combine.PartWhole (part-part-whole aggregation) / Combine.Product (structured product:  $m \times n$ , area/arrays, combinations)
- Compare.Difference (more/less, difference) / Compare.Ratio (times-as-many, each/per, rates)
- Equalize.MatchDifference (add/remove to match) / Equalize.MatchRatio (scale to match)
- Other

#### B) Arithmetic operations (use when building equations)

1. ADDITION (+): combine distinct quantities to find a total (sum, total, altogether, receive).
2. SUBTRACTION (-): difference, removing a quantity, or additive comparison (left, fewer, lost, give away).
3. MULTIPLICATION (\*): repeated groups, scaling, or structured product (times, each, double, rows  $\times$  columns).
4. DIVISION (/): splitting into equal parts or solving inverse of a multiplicative relation (share, split, average, per-unit).

Always ground equations in the entities and relations described in the problem.

#### DEMONSTRATIONS:

**Q:** There are 15 trees in the grove. Grove workers will plant trees in the grove today. After they are done, there will be 21 trees. How many trees did the grove workers plant today?

#### Discourse Information:

1. "There are 15 trees in the grove."
2. "Grove workers will plant trees in the grove today." (Relation: Change.Increase)
3. "After they are done, there will be 21 trees." (Relation: Change.Increase Result)
4. "How many trees did the grove workers plant today?"

**A:** Math-DB: The situation is a Change relation: planting increases the number of trees, and we are given the final result.

Let start = 15 and result = 21. This is Change.Increase with unknown\_role = delta (how many were added). So delta = result - start = 21 - 15 = 6. The final answer is 6.

**Q:** If there are 3 cars in the parking lot and 2 more cars arrive, how many cars are in the parking lot?

#### Discourse Information:

[...Discourse Clauses and Relations...]

**A:** Let start = 3 and delta = 2, with unknown\_role = result. So result = start + delta = 3 + 2 = 5. The final answer is 5.

[...Remaining Demonstrations Omitted...]

**Q:** {target\_question}

**Discourse Information:**

{target\_discourse\_info}

**A:** Let's think step by step.

### Prompt Template: Baseline CoT (GSM-Symbolic)

// preamble or system instruction

As an expert problem solver, solve step by step the following mathematical questions.

// shot-1

Q: There are 15 trees in the grove. Grove workers will plant trees in the grove today. After they are done, there will be 21 trees. How many trees did the grove workers plant today?

A: Let's think step by step. There are 15 trees originally. Then there were 21 trees after some more were planted. So there must have been  $21 - 15 = 6$ . The final answer is 6.

.  
.  
.

// shot 8

Q: Olivia has \$23. She bought five bagels for \$3 each. How much money does she have left?

A: Let's think step by step. Olivia had 23 dollars. 5 bagels for 3 dollars each will be  $5 * 3 = 15$  dollars. So she has  $23 - 15$  dollars left.  $23 - 15$  is 8. The final answer is 8.

// target question

Q: {{{question}}}

A: Let's think step by step.

## D Appendix: Inter-Annotator Agreement

Although we used majority vote rather than measuring exact agreement, we did compute inter-annotator agreement on a randomly sampled subset of 1,000 relations to measure the difficulty of the annotation task. Among the three annotators on each relation, the initial agreement (before adjudication), measured by Fleiss'  $\kappa$ , was 0.84 on the Level-2 sense. This is reasonably high, indicating that the framework was clear for most cases. Fleiss'  $\kappa$  for the five Level-1 sense categories was 0.92, showing substantial agreement. The most confusion was in differentiating Change vs Combine in cases of additive actions, and identifying Other vs a weak Combine/Change link. After discussions and adjudication, those were resolved in the final version of the Math-DB.

Overall, the Math-DB annotation process yielded a high-quality dataset of discourse-labeled math problems. This resource not only enables our experiments but can also help future studies, such as training automated discourse parsers for math text or analyzing common patterns in math problem structure.

# Cross-Linguistic Situation Entity Segmentation for Discourse Analysis in Diachronic English and German Text

Hanna Schmück, Veronika Urban, Xaver Krückl, Sonja Zeman,  
Claudia Claridge, and Annemarie Friedrich  
University of Augsburg, Germany

## Abstract

Situation Entity (SE) segmentation identifies clause-like discourse units focusing on verb constellations. While SE segmentation has been applied to contemporary English as a sub-task of SE annotation, systematic guidelines for syntactically ambiguous constructions remain underspecified. We present principled SE segmentation guidelines for contemporary and historical varieties of English and German. Our inter-annotator agreement studies on Late Modern English (1700–1900) and New High German (1650–1900) corpora demonstrate substantial agreement. Using the existing SitEnt corpus in contemporary English, we implement a new automatic segmenter based on XLM-RoBERTa. Our evaluation examines cross-variety and cross-lingual generalization, demonstrating challenges both for human annotation efforts and in transferring segmenters trained on contemporary English to historical varieties. Our code and data are publicly available.<sup>1</sup>

## 1 Introduction

Discourse segmentation, the task of identifying meaningful units of text, is fundamental to computational discourse analysis. One framework for discourse analysis and segmentation are Smith’s (2003) modes of discourse. These modes result from specific situation entity (SE) patterns, and SE segments are the focus of our annotation task. This framework is of particular interest since it offers a linguistically motivated approach of segmenting texts which is grounded in aspectual properties of clause-like elements. An example is visualized in Figure 1. SE segmentation serves as a prerequisite for two core downstream tasks: Firstly, it can be used as a basis for discourse mode classification, a valuable tool for the analysis of literary and historical corpora where shifts between, e.g., Narrative

<sup>1</sup><https://github.com/coling-unia/sitent-segmenter-law2026>



Figure 1: Situation annotation of New High German (top, Koralek (1889/1890)) and Late Modern English (bottom, Salmon (1724)) text snippets showing one segment per line and the respective Main Verb and Main Referent relations. Main Verbs can be connected to their respective Main Referent across segment borders.

and Information mode reflect authorial style, genre conventions, and rhetorical structure. Secondly, SE segments correspond roughly to propositions and therefore constitute a linguistically motivated unit for information and event extraction in historical texts.

Manual segmentation of text into clauses or clause-like units is not a trivial task, and requires the careful design of annotation guidelines. It has been approached both from a syntactic perspective (Bies et al., 1995) and a discourse perspective (Polanyi, 1995; Carlson et al., 2001; Polanyi et al., 2004). Automatic clause segmentation is also non-trivial (Tjong et al., 2001; Soricut and Marcu, 2003; Tofiloski et al., 2009). As Carlson et al. (2001) put it, “the boundary between discourse and syntax can be very blurry.”

While SE annotation has been explored for contemporary German (Mavridou et al., 2015) and applied to English (Friedrich et al., 2016), existing guidelines for segmentation (Friedrich et al., 2015a, chapter 3) mainly specify how annotators are supposed to handle differences vs. the EDU segmentation produced by a legacy system (Soricut and Marcu, 2003).

In this work, we address these challenges by providing more systematic SE segmentation guidelines for both contemporary and historical varieties of English and German. We thereby close the gap left by previously undocumented approaches and generalize the idea to new language variants. We further conduct a small IAA study on SE segmentation of these varieties based on the Corpus of Late Modern English Historiography (CLMEH (Claridge, 2025)) and the GiesKaNe corpus (Justus-Liebig-Universität Gießen, 2022). Lastly, we fine-tune XLM-RoBERTa models for automatic SE segmentation and evaluate their cross-variety generalization abilities. Our automatic segmenter achieves a boundary F1 of 0.90 on contemporary English, and human annotators maintain substantial agreement across all four varieties (75–84% exact matches), demonstrating that the guidelines can be adapted to suit cross-lingual and historical settings.

## 2 Background

We provide an overview of SE types and discuss the broader research area of discourse segmentation.

### 2.1 Linguistic Background

*Discourse modes* (Smith, 2003) refer to different text types (Werlich, 1989; Biber, 1989; Adam, 2011), which are characterized by clusters of linguistic features. In Smith’s *Narrative* and *Report* modes, for example, the reader perceived the discourse to move on predominantly via temporal relations. By contrast, progression is perceived as spatial in *Description* mode, like traversing through a scene. In *Information* and *Argument* mode, the discourse advances by focusing on different referents that are part of the domain of the discourse.

Discourse modes are not only distinguished by their type of progression: they also favor different distributions of *SE types*. SE types capture the aspectual form of (roughly) clauses. We illustrate and describe the inventory of SE types as proposed in prior work (Friedrich and Palmer, 2014b; Friedrich et al., 2016) in Table 1.

According to Smith (2003), SE types are assigned to *verb constellations*, i.e., a verb and its “primary referent” in the context of its arguments and modifiers. A crucial step in SE segmentation hence lies in determining what counts as a *main verb*. Friedrich et al. (2016) flesh out three features that help to distinguish SE types for new annotators or to obtain partial information if not all features

SE type	Example
<b>Eventualities</b>	
STATE	The colonel owns the farm.
EVENT	John won the race.
REPORT	“...”, said Obama.
<b>General Statives</b>	
GENERIC SENT.	The lion has a bushy tail.
GENERALIZING SENT.	Mary often fed the cat last year.
<b>Abstract Entities</b>	
FACT	I know that she refused the offer.
PROPOSITION	I believe that she refused the offer.
QUESTION	Who wants to come?
IMPERATIVE	Hand me the pen!

Table 1: SE type inventory (Smith, 2003; Friedrich et al., 2016; Friedrich and Palmer, 2014b).

can be determined. First, annotators determine whether the *main referent*, in English typically the subject of the main verb, is generic (Krifka et al., 1995) or whether it refers to a particular individual (see also Nedoluzhko, 2013; Friedrich and Pinkal, 2015b; Friedrich et al., 2015b). Generic main referents identify GENERIC SENTENCES. Second, the lexical aspectual class (Moens and Steedman, 1988; Klavans and Chodorow, 1992; Friedrich and Palmer, 2014a) plays a role, e.g., for distinguishing EVENTS from STATES. Third, if the main verb is *habitual*, i.e., if it indicates a situation that happens regularly, the SE must be either a GENERIC or a GENERALIZING SENTENCE (Vendler, 1957; Mathew and Katz, 2009; Friedrich and Pinkal, 2015a). The full correspondence between the lower-level aspectual features and SE types are provided in Appendix B.

Challenges regarding SE boundary identification intensify when working with historical language varieties. Late Modern English (LModE; approximately 1700–1900) and New High German (NHG; approximately 1650–1900) exhibit syntactic features that complicate SE boundary identification. LModE employs participial absolutes, reduced relative clauses, and infinitival purpose constructions with less explicit subordination marking than contemporary varieties. NHG shows a less restricted word order, particularly in verb positioning, when compared to present day German and makes use of participial clauses and final clause constructions with *damit* which introduce new boundary cases.

## 2.2 Existing Discourse Segmentation Guidelines

The first step when parsing text into discourse structures consists in determining which linguistic units correspond to semantic units in the discourse. Several frameworks have proposed different solutions to this problem. The Linguistic Discourse Model (Polanyi, 1995; Polanyi et al., 2004) defines Basic Discourse Units (BDUs) as segments with the potential to establish anchor points for future attachment, identifying syntactic constructions able to carry the necessary semantic information. In practice, BDUs are often rather small units, e.g. in the following case:

(1) [Germany elected Merz] [chancellor].

This is more fine-grained than SE annotation, since *chancellor* does not constitute a separate situation.

The Penn Discourse Treebank (PDTB) (Mitsakaki et al., 2004), distinguishes inter-sentential and intra-sentential discourse relations between segments called *argument spans*. The latter relate subordinating clauses, complement clauses, free to-infinitives, and nominalizations as well (Webber et al., 2019), as illustrated in examples (2)-(4).

(2) [Treasurys opened lower] ARG1, [reacting negatively to news]ARG2

(3) ... [a number of project veterans were on hand to watch the launch]ARG2 [to watch the launch]ARG1

(4) ... many are hoping [for major new liberalizations]ARG2 if [he is returned firmly to power.]ARG1

Rhetorical Structure Theory (RST, Mann et al., 1992; Carlson et al., 2001), on the other hand, defines Elementary Discourse Units (EDUs) essentially as clauses. Their size is in principle arbitrary, but the units should have independent functional integrity (Mann and Thompson, 1988). Clausal subjects, complements and restrictive relative clauses are considered as parts of the clause headed by their governing verb. Applying this intuition consistently at scale requires extensive rule sets motivated by the inventory of discourse relations. For more extended examples of EDUs as well as a Figure comparing the alternative segmentation approaches discussed in this section with SE segmentation see Appendix A.

## 2.3 Automatic Discourse Segmentation

The first steps in neural approaches to automatic discourse segmentation were taken by Wang et al.

(2018) in NeuralEDUSeg, applying a BiLSTM-CRF architecture and addressing data sparsity through pre-trained word embeddings plus a restricted self-attention mechanism. In the context of EDU segmentation, the DISRPT shared tasks have driven further progress on neural approaches. The 2019 winning system, ToNy (Muller et al., 2019), combined a BiLSTM-CRF with multilingual BERT-based sequence prediction across 15 corpora spanning RST, SDRT and PDTB, outperforming prior models on nearly all languages. Next, Gessler et al. (2021) extended ToNy by incorporating token-level handcrafted features such as POS tags and dependency relations. More recent winning systems at DISRPT 2023 by (Braud et al., 2023) and 2025 by (Lalitha Devi et al., 2025) build on using XLM-RoBERTa (Conneau et al., 2020). Braud et al. (2023) additionally experiment with freezing specific layers to separate morpho-syntactic from semantic encoding. As a contrast, using a generative language model, Nayak (2024) showed that zero-shot prompting of GPT-3.5 turbo is still not competitive with smaller pretrained language models explicitly trained for segmentation. Frenzel et al. (2026) provide the most recent overview on the development from rule-based systems to neural approaches, and similarly adopt fine-tuning XLM-RoBERTa as the segmentation model in their approach on German data.

Previous methods for automatic SE segmentation (Friedrich, 2017) add a postprocessing step to SPADE to convert RST EDUs to SE segments. The segmenter of the discourse parser SPADE (Soricut and Marcu, 2003) is based on a probabilistic model learned from the RST Discourse Treebank. For all words in the vocabulary, the probability for inserting a boundary after a word  $w$  is estimated from the treebank and depends on a lexicalized version of the corresponding sentence’s syntactic tree (Magerman, 1995).

## 3 Segmentation Guidelines for Situation Entities

Smith (2003) suggests that SEs are introduced by the clauses of a text, while noun phrases introduce individuals (e.g., people, places, objects or ideas) and tense and time adverbs introduce times. She does not further specify which linguistic construction she considers to be a clause. Following the idea that discourse segmentation depends on linguistic units to which the target categories can meaning-

fully be assigned, in this work, we define and operationalize a segmentation scheme for SE types for English and German.

### 3.1 Situation Segments

We now describe which syntactic units can be assigned SE types in English, i.e., which constructions we assume to function as verb constellations.

**Finite clauses** such as Example (5) are the most clear-cut case. The finite verb is the SE’s main verb, its grammatical subject is the main referent. Earlier work on SEs (Friedrich et al., 2016) is based on Stanford dependencies (de Marneffe and Manning, 2008), which mark participles as dependents of the finite auxiliaries or modal verbs. In the more recent Universal Dependencies framework (de Marneffe et al., 2021), auxiliaries are dependents of the participles because those carry the semantic meaning of the verbs. In both the earlier work and our recent extensions, the main verbs are marked on the same spans as the participle.

- (5) John<sub>MREF</sub> built<sub>MV</sub> a house.  
Whales<sub>MREF</sub> are<sub>MV</sub> huge.

When the finite verb is an auxiliary, the main verb for the purpose of SE annotation is still the meaning-carrying verb even if this is non-finite (see Example (6)).

- (6) He<sub>MREF</sub> may join<sub>MV</sub> us later.  
This<sub>MREF</sub> won’t help<sub>MV</sub>.

The segmentation rule also applies to expletive sentences where “it” functions as an empty subject.

- (7) It is raining.  
 It turned dark.

While there is no clear main referent in these cases, it is typically still possible to decide whether a statement is made about a class or kind (generic) or not, as in the segments listed in (7).

**Relative clauses** clearly introduce SEs.

- (8) (a) [My brother<sub>1</sub>], [who lives in Chicago<sub>2</sub>],  
 [is visiting this weekend<sub>1</sub>].  
 (b) [The book<sub>1</sub>] [that fell down<sub>2</sub>] [was mine<sub>1</sub>].

As illustrated in example (8), SE segments are not necessarily contiguous spans of text, which makes them stand out compared to the discourse segmentation schemes discussed in Section 2.2. For practical reasons, we split the complete text into segments, but not each of the segments will ultimately be assigned SE features. The noun phrases “My brother” and “The book” will simply be labeled with NO-SE, however, they syntactically function as the subject and as the main referent

of the matrix clauses. Annotators are instructed to consider the underlying syntactic dependencies when deciding on SE labels. We opted for this approach due to its simplicity during annotation and also to emphasize that SE types capture the form of linguistic realizations rather than more abstract semantic notions as, e.g., in Abstract or Unified Meaning Representations (Banarescu et al., 2013; Gysel et al., 2021).

**Participles** are not considered to evoke SEs if their use is purely adjectival as in “The dancing girl.” The reasoning behind this is that we aim for a granularity that roughly corresponds to clauses. By contrast, participle clauses are also considered to be segments, e.g., if they function as reduced relative clauses (9) or participle clauses indicating temporal (10) or causal (11) relations. As a general rule, proposed participles without modifiers or compliments tend to be read as adjectival and therefore do not constitute their own segment whereas postponed participles are rarely read in purely adjectival manner and therefore invoke a separate situation.

- (9) (a) [The man] [talking to John] [is my brother].  
 (b) [The book] [written by Orwell] [is famous].  
 (10) (a) [Walking home], [I met Sarah].  
 (b) [Having finished], [I got up].  
 (11) (a) [The bomb exploded], [destroying the bridge].  
 (b) [Loving her], [he proposed].

Our decision to treat participle clauses but not participles in adjectival use as SEs is rooted in the idea that they are more similar to the other types of clauses we define. We acknowledge, however, that this is not a clear-cut case.

**Gerunds** that are used as nouns, e.g., “running” or “walking” in the full segment (12) are, analogously to the BDU approach, not considered to invoke their own SEs. These gerunds clearly refer to concepts instead of situations.

- (12) Running burns more calories than walking.

Asher (1993) investigates how eventualities (states and events) and abstract entities (propositions, properties, states of affairs and facts) are referred to in natural language. He provides an inventory of sentential nominals, i.e., syntactic structures whose meanings are correlated with sentences: derived nominals (13), gerund phrases (14), that-clauses (15), for-infinitival phrases (16), naked infinitive phrases (3a) and noun phrases involving common nouns that may combine with that-clauses

or gerund phrases (17).

- (13) (a) The army’s destruction of the city
- (b) Franklin’s favorite invention
- (14) (a) The mayor’s throwing of the pizza
- (b) John’s hitting Bill
- (c) The gathering of the pecans
- (15) that Sam greeted Susan
- (16) (a) For Kim to win was unexpected.
- (b) John wanted for Mary to be chair.
- (17) (a) Mary’s doubt that John was unhappy
- (b) The fact that John was unhappy
- (c) The letter explaining the situation

Some of these constructions such as (13a) or (15) clearly refer to events while others may also refer to objects (13b). Of the above sentential nominals, we mark only that-clauses as SE segments. We do not treat the rest of the above sentential nominals as invoking SEs, but simply as part of the larger situation segment in which they are embedded. We decided not to mark them because the boundary between event-denoting constructions (13a) on the one hand and phrases denoting concrete (13b) or abstract objects (17) on the other hand is not clear; additional annotation guidelines would be necessary.

**Conjunctions and conditionals.** When a clause starts with a conjunction or a subordination, the conjunction or subordination is segmented into the same span as the clause that they introduce.

- (18) (a) [I hate] [and love him].
- (b) [I believe] [that she called him].
- (c) [I left] [after I had called him].
- (d) [She left] [because I had called him].
- (e) [I like to sit] [where the sun shines].
- (f) [I like to run] [if the sun shines].

Please note that our general segmentation depends on whether aspectual features can be assigned to the verb construction expressed by a span of text. It does not matter which role the sub-clause plays in the argument/modifier structure of the embedding verb, e.g., in (18b), the embedded segment is a complement clause, while (18c) and (18d) are adjunct clauses.

**To-infinitives** are not considered to introduce SE segments (Friedrich et al., 2015b). In English, to-infinitives can fulfill different functions. In (19a) and (19b), the to-infinitives are predicative nominalizations and refer to abstract concepts. To-infinitives can also indicate purpose (19c) or function as adjectival or other complements as in (19d) and (19e). By definition, we assume these cases to be predicative uses rather segments that introduce

their own situation entities.

- (19) (a) [To travel alone can be exciting.]
- (b) [My dream is to study abroad.]
- (c) [She went to the library to study.]
- (d) [They are happy to be here.]
- (e) [He grew up to become a teacher.]

The core idea here is that SEs annotate the aspectual forms as chosen by the writer. While this does require semantic interpretations, we aim not to reformulate too much, as this introduces too many degrees of freedom.<sup>2</sup>

### 3.2 Extension to German

In historical language data, sentence boundaries cannot always be identified based on punctuation, as this differs significantly from that of modern English and German. In the case of the GiesKaNe corpus, this problem was already resolved during corpus creation, and sentence-level annotation, in accordance with the annotation guidelines, does not take place at the orthographic level. Sentences are thus determined grammatically, not by punctuation (Ágel and Henning, 2023). German syntax from the 17th to the 19th century also poses particular challenges for SE segmentation. The following subsections describe the resulting specific adjustments to the annotation scheme with regard to two selected phenomena. As work continues with additional data, further phenomena may be identified.

**Participle clauses.** In contemporary German, as in NHG, participles can be used in attributive function, in which case they are not classified as segments. In contrast to contemporary German, it is much more common in NHG to have participial clauses where the finite auxiliary is missing, also known as "auxiliary ellipsis" (Breitbarth, 2005; Thomas, 2018). This is actually a phenomenon of Early NHG, but its residual effects are still relevant for NHG data. Participial clauses without finite auxiliaries can be classified as segments, as has already been argued for English. In many cases, the heuristics from English do not work due to the less restricted word order (Ágel, 2015, 2000) and/or the missing finite verb in NHG. Even if parts of sentences are recognizable as reduced relative clauses (20a) or as clauses with auxiliary ellipsis (20b), only in some cases does the corresponding word order allow for segmentation in the NHG data.

<sup>2</sup>A special case where we allow minor hypothetical reformulations are participle clauses, e.g. *Sitting on the bench, I looked at the beach.* – *While I was sitting on the bench, I looked at the beach.*

- (20)(a) [So sieht man denn selbst so genannte Gebildete, — subjektiv stumpf, objektiv peinigend — ... die lieblichsten, [durch die Natur vorgebildeten,] [durch die Kochkunst veredelten Produkte] naturalistisch und roh sich aneignen].<sup>3</sup> (Anthus, 1838)
- (b) [Er erschiene unsäumig;] [und als Sie sich mit ihm an ein Fenster gesteuert] / [fragte Sie]<sup>4</sup> (Birken, 1652)

**Final clause construction.** To mark final clause constructions in German, we follow Mavridou et al. (2015). German provides the construction of final clauses with “damit” or “so ... dass” (“such/so that”), indicating a purpose or goal (21a) or an actual event (21b).

- (21)(a) [Erinnere mich nochmal,] [damit ich pünktlich komme.]<sup>5</sup>
- (b) [Da stopft Einer gedankenlos ... so viel Brod in den Mund,] [daß er unmöglich den spezifischen Geschmack irgend einer Speise perzipieren kann.]<sup>6</sup> (Anthus, 1838)

As the subordinated clauses are final, we treat them as separate segments containing the conjunctions introducing the clause.

### 3.3 Extension to LModE

In addition to covering peculiarities of German with regards to segmentation we further provide a brief description of a special case arising when annotating LModE.

**Absolute constructions.** One significant difference between LModE and Present Day English is the higher frequency of use of the absolute construction (Van De Pol and Petré, 2015), a common boundary case. Example (22) begins with such a construction.

- (22) [The Parliament being met on the 23d October,] [his Majesty refer'd them to what he had said to both Houses four Days before.] (Boyer, 1702)

These constructions present their own SEs without finite verb morphology or subordinating conjunctions. Their syntactic independence suggests

<sup>3</sup>English: [Thus one sees even so-called educated people, — subjectively dull, objectively tormenting — ... appropriate the loveliest products, [pre-formed by nature,] [refined by the culinary arts ] naturalistically and crudely.] Note that in the original, the descriptions of the products are preposed.

<sup>4</sup>English: [He appeared without delay;] [and when steered herself with him to a window] / [she asked]

<sup>5</sup>English: [Remind me again] [so that I arrive on time.]

<sup>6</sup>English: [There someone thoughtlessly stuffs so much bread into their mouth with every bite] [that they cannot possibly perceive the specific taste of any dish.]

separate segmentation, yet they can function as adverbial modifiers of the main clause. The lack of explicit conjunctions makes their relationship to the main clause SE less transparent than in equivalent finite subordinate clauses (*When the parliament had met...*). These purpose constructions occupy an intermediate position between arguments and adjuncts. They describe potential or foundational situations rather than actualized events, yet lack the full clausal structure of *that*-complements or *so that* result clauses. Despite this and due to their syntactic independence they are marked as separate segments in this project.

## 4 Annotation Study

In order to extract relevant underspecified edge cases, and to evaluate IAA as well as model performance, we conducted a human annotation study spanning four language varieties, contemporary and historical English and German, using excerpts from written corpora.

### 4.1 Data Sources

Four datasets are used for this annotation study, corpus statistics including the number of annotated tokens per variety can be found in Table 2.

The Corpus of English Historiography (CLMEH (Claridge, 2025)) is used as the basis for segmentation of historical English text. For the manual annotation, two snippets from two different LModE authors, Abel Boyer and Thomas Salmon, were annotated. These texts were chosen since they are among the earliest texts in the corpus (published in 1702 and 1736 respectively), and therefore serve as a stress-test of the annotation guidelines.

The NHG data is a subset of the GiesKaNe Korpus (Justus-Liebig-Universität Gießen, 2022), similarly consisting of early texts by Johann Joachim Becher (1668) and Antonius Anthus (1838).

For modern German we use snippets from German Wikipedia articles as well as four blog articles downloaded from publicly available and CC-BY-SA-licensed blogs as listed in DWDS (Barbaresi and Würzner, 2014).

Lastly, to estimate human agreement for contemporary English, we also annotate present-day English based on a subset of four files from the the MASC+wiki dataset used in Friedrich (2017), one travel blog, a news story, an email, and an English Wikipedia page. Appendix C contains information on the text sources including author names where

available, as well as the year of publication for each manually annotated document, Table 2 contains the token counts per variety.

Language Variety	# Tokens	# Segs	Avg. Tok/Seg	Max. Tok/Seg
Current E	3,937	482	10.3	42
LModE	6,797	614	13.6	56
Current G	4,751	444	12.6	64
NHG	5,666	617	11.4	67

Table 2: Volume of annotated data per language variety.

## 4.2 Inter-Annotator Agreement

We collect two independent human segmentation annotations per language variety, which we call A and B<sup>7</sup>. The human segmentation was performed by five native German speakers as annotators, three authors of this paper and two paid student annotators. All annotators underwent several weeks of specific and language-/variety-dependent segmentation training. We evaluate annotation reliability using exact span match agreement, computed as the proportion of segments in a reference annotation that are recovered with identical span boundaries in the comparison annotation. Annotators were trained jointly but annotated independently. Results can be found in Table 3.

Agreement is computed over span boundaries derived from pre-tokenized BI representations using the same tokenizer as for model training (see Section 5). As the task is unitizing (segmentation without labels) over continuous text, standard chance-corrected measures such as Krippendorff’s  $\alpha_U$  are not applicable (Artstein and Poesio, 2008). The left column of Table 3 provides annotation statistics and agreement scores across all four varieties. Overall average agreement reaches 79.1%, with all varieties falling within a narrow 5-point range, supporting the general robustness and cross-linguistic transferability of the guidelines. Contemporary English and Contemporary German show the largest exact span match percentages – on average 81.4% and 80.2% respectively. As expected, agreement for historical variants is harder to achieve and sits at 76.86% on average for NHG and 77.98% for LModE.

To gain a better understanding of the root causes of all disagreements, they were explored by one of

<sup>7</sup>Please note that the A and B labels do not always refer to the same human annotators

the authors and manually assigned a disagreement group and label.

A closer look at the inter-annotator disagreement (see Figure 3 in Appendix C for an exact breakdown disagreements) reveals that inconsistencies concentrate in syntactic and discourse-related phenomena. Across all four conditions, *clause missed* — scenarios where a clause containing a valid meaning-carrying verb was disregarded — constitutes the single most frequent source of disagreement. As explored above, the NHG condition narrowly shows the largest overall disagreement counts, with particularly high occurrences of *clause missed* disagreements and its sub-category of missed *relative clauses*. We attribute this to phenomena such as frequent (auxiliary) ellipses, long sentences containing multiple clause chains that are unusual to the modern annotator, and multiple valid boundary placement options due to a less restricted word order.<sup>8</sup> Current German disagreements cluster around *inlay* and *arbitrary boundary* placement since word order is less restricted leading to multiple valid splitting points within long sentences. Current English exhibits elevated disagreement for *verbless* and *participle* constructions alongside *arbitrary boundary* decisions, pointing to uncertainty in identifying implicit or non-canonical clause structures. LModE shows a similar profile, with persistent disagreement in *clause missed* and *verbless* constructions, as well as the discourse-adjacent category of *absolute constructions*<sup>9</sup>. Orthographic disagreements remain rare across all datasets. These findings suggest that disagreement primarily arises from complex syntactic boundary decisions and discourse-level interpretation, and is amplified in historical data where structurally ambiguous constructions are more frequent.

<sup>8</sup>See this excerpt from Anthus as an Example: *Wie aber die Gelehrten noch dar über schwitzen , zu bestimmen , wo die Pflanze zu m Thier wird , wo die Grenzen des Pflanzen- und Thierreiches fest zu stecken seien , eben so schwierig ist es , zu bestimmen , welcher unteren Thierreihe man zuerst die Fähigkeit eines eigentlichen Essens zu zu gestehen hat.* English: *But as the scholars still sweat over how to determine, where the plant becomes the animal, where the borders of the plant and animal kingdoms are to be placed, as hard is it, to determine, to whom amongst the order of animals one should first grant the property of a true food.*

<sup>9</sup>See this excerpt from Boyer (LModE) as an example: *On the 29th of October 1689, came on the usual Solemnity of the Lord Mayor of London, and Sir Thomas Pilkinton being continued for the Year 1690, and the King and Queen, the Prince and Princess of Denmark, and both Houses of Parliament, having been pleas’d to accept his Invitation to his Dinner, their Majesties attended by their Royal Highnesses, [...].*

Variety	Inter-annotator agreement (exact match)		Model performance (exact match)	
	Ann. B vs. A	Ann. A vs. B	Model vs. Ann. A	Model vs. Ann. B
Contemporary G	76.48±14.25	83.92±11.83	69.32±12.22	64.20± 6.73
NHG	78.36± 0.75	75.36±11.56	42.68±30.51	48.87±38.06
Contemporary E*	80.62±16.69	82.15±12.25	77.06± 8.58	86.91± 9.36
LModE	79.34± 1.04	76.62± 4.88	64.28± 3.21	67.50± 2.90

Table 3: Inter-annotator exact span match (%) per variety. Agreement is reported once with annotator A treated as the reference (gold) annotation and once with annotator B treated as reference, i.e., the first column shows how many of A’s annotations were exactly matched by B and the second column vice versa. Model performance as exact span match (%) versus human annotator A and B. \*As the model’s training set contains the test documents for Contemporary English, we trained a separate model for this setting without these test documents in the training data.

## 5 Automatic Segmentation Model

We implement a new SE segmentation model that does not rely on legacy systems, using about 40,000 labelled English SE segments provided by Friedrich et al. (2016).<sup>10</sup> The training data includes texts from a variety of genres taken from the Manually Annotated Subcorpus of the OANC (MASC) (Ide et al., 2008) and Wikipedia. The original rule-based segmentation model was based on a legacy system (Soricut and Marcu, 2003) which itself had been trained on only just over 7,000 sentences. While we do not yet have sufficient training data for German and historic English variants, we apply our model to these variants in a zero-shot way, achieving promising results.

### 5.1 Modeling

XLm-RoBERTa (Conneau et al., 2020) has proven itself a robust baseline for EDU segmentation not only for English (Braud et al., 2023; Lalitha Devi et al., 2025), but also for German data (Frenzel et al., 2026). Thus, we similarly treat segmentation as a sequence tagging task and fine-tune XLm-RoBERTa (Conneau et al., 2020), a multilingual transformer-based encoder model producing contextualized token representations. We preserve the predefined train-test splits on document level and sample a development set from the training split across documents, resulting in a final train/dev/test ratio of 0.68/0.10/0.22. First, we segment each document into sentences using spaCy’s (Explosion AI, 2025) en\_core\_web\_sm pipeline and treat each sentence as an independent input.<sup>11</sup> To mark segment boundaries, we use

<sup>10</sup><https://github.com/annefried/sitent>

<sup>11</sup>This modular design ensures that, in future work, we can train sentence segmenters using any sentence-segmented labeled data, which is important particularly for the historical language variants, where training data is limited and conventional tokenization may fail, allowing to fine-tune segmenters

Metric	Avg.	SD
B F1	90.5	0.1
B Precision	90.4	0.7
B Recall	90.6	0.8
Exact Match	74.3	0.3
WindowDiff	7.7	0.1

Table 4: Situation segmentation performance (detecting start boundary tokens) on the SitEnt test set using XLm-RoBERTa-large, averaged over 5 random seeds. This test dataset had been single-annotated by Friedrich et al. (2016); the type of data is most similar to Contemporary E, where human agreement reached up to 82.15 % exact match.

a BI sequence tagging scheme. Following common practice, we backpropagate the loss for each first SentencePiece (Kudo and Richardson, 2018) token corresponding to a spaCy token. We fine-tune XLm-RoBERTa-large with a linear classification head and a Conditional Random Field (CRF) (Lafferty et al., 2001) decoding layer.

### 5.2 Evaluation Metrics

Next to precision, recall and F1 score, we also report exact matches and WindowDiff (Pevzner and Hearst, 2002) using the NLTK implementation (Bird and Loper, 2004). WindowDiff evaluates segmentation quality by sliding a fixed-size window over both the reference and predicted boundary sequences and comparing the number of boundaries within each window. Instead of requiring exact boundary matches, it penalizes differences in local boundary counts. See Appendix D for details on hyperparameters and training procedure.

### 5.3 Experimental Results

Results for XLm-RoBERTa-large on the SitEnt test set are reported in Table 4. The averaged results over five random seed runs show a well- on smaller, sentence-aligned datasets.

balanced trade-off between precision and recall, with a slight tendency of the model to overpredict segments. While Exact Match is considerably lower than the boundary detection scores, the low WindowDiff suggests that most errors stem from minor boundary misalignments (e.g., slightly shifted EDU boundaries or punctuation handling), rather than complete segmentation failures.

In the next step, we apply our best performing model on boundary F1 across the hyperparameter grid to annotate the datasets described in Section 4. Table 3 reports exact matches between the model output and each of the two independent human annotations per language variety. Exact match is again calculated as in the IAA evaluation above. For each document, the model annotation is compared independently against each of the two human annotations, yielding two exact match scores per document.

#### 5.4 Discussion of Computational Experiments

The results in Table 3 reflect the expected degradation in model performance the more the language variety diverges from the model’s training domain of contemporary English. On Contemporary E, the model reaches mean exact match scores of 77.06% (vs. Ann. A) and 86.91% (vs. Ann. B), approaching human agreement levels (IAA = 81.39%), proving that the model captures segmentation patterns in in-domain data reasonably well. It is much closer to Ann. B who has been part of the training data annotation. Applying the model cross-lingually to Contemporary G reveals a substantial drop to 69.32/64.20%, despite human IAA remaining stable at 83.92%.

While human annotators remain highly consistent in German, the model fails to fully transfer its segmentation knowledge across languages. For the historical varieties a similar pattern of degradation becomes apparent. For LModE, the model achieves 64.28%/67.50%, compared to a human IAA of 79.34%/76.62%, reflecting a certain difficulty with historical English. A low cross-document variability ( $SD = 3.21/2.90$ ) suggests that this difficulty is consistent across texts. Contemporary G and LModE show highly similar model performance ranges, suggesting that cross-lingual variation in contemporary German and diachronic variation in Late Modern English impose a comparable level of difficulty for the model. NHG represents the most challenging setting for the model, combining cross-lingual and temporal domain shift, leading to the

lowest exact match at 42.68/48.87%, against a human IAA of 78.36%/75.36. A high cross-document variability ( $SD = 30.51/38.06$ ) indicates a substantial variation across the annotation of the two evaluated texts. While human annotators remain relatively consistent, the model fails to recover more than half of the same boundaries under this combined condition.

## 6 Conclusion

We have presented principled SE segmentation guidelines for contemporary and historical varieties of English and German, filling a gap left by prior work which relied largely on implicit or underspecified segmentation criteria. Our inter-annotator agreement study demonstrates substantial and consistent human agreement across all language varieties, with average exact match scores ranging between 76.86% and 81.4%. This confirms that the guidelines are, in principle, learnable and robust even for syntactically challenging historical texts. Our new XLM-RoBERTa-based segmenter achieves a boundary F1 of 90.5 on contemporary English, matching human agreement on its training domain. However, model performance drops substantially under cross-lingual and temporal domain shift, most severely for historical German, where exact match falls to 45.78%. Although the segmentation task itself is consistent across varieties, the bottleneck lies in the availability of in-domain training data. We release our guidelines, annotated data, and model code to support future work on multilingual and diachronic discourse analysis.

## Acknowledgement

The authors gratefully acknowledge the scientific support and HPC resources provided by the Erlangen National High Performance Computing Center (NHR@FAU) of the Friedrich-Alexander-Universität Erlangen-Nürnberg (FAU) under the BayernKI project v110ee. BayernKI funding is provided by Bavarian state authorities. We thank Manfred Pinkal and Alexis Palmer for their helpful discussion related to this annotation scheme. We also extend our thanks to the annotators as well as to the anonymous reviewers for their constructive and useful feedback.

## References

- Jean-Michel Adam. 2011. *Les textes: types et prototypes: récit, description, argumentation, explication et dialogue*. Armand Colin.
- Vilmos Ágel. 2000. *Syntax des neuhochdeutschen bis zur mitte des 20. jahrhunderts*. In Werner Besch, Anne Betten, Oskar Reichmann, and Stefan Sonderegger, editors, *Sprachgeschichte, Part 2*. Walter de Gruyter, Berlin, New York.
- Vilmos Ágel. 2015. *Die umparametrisierung der grammatik durch literalisierung. online- und offlinesyntax in gegenwart und geschichte*. In Ludwig Eichinger, editor, *Sprachwissenschaft im Fokus*, pages 121–156. DE GRUYTER.
- Vilmos Ágel and Mathilde Henning. 2023. *Annotationshandbuch des dfg-projekts syntaktische grundstrukturen des neuhochdeutschen. zur grammatischen fundierung eines referenzkorpus neuhochdeutsch*.
- Antonius Anthon. 1838. *Vorlesungen über Esskunst*. Wigand, Leipzig.
- Ron Artstein and Massimo Poesio. 2008. *Inter-Coder Agreement for Computational Linguistics*. *Computational Linguistics*, 34(4):555–596.
- Nicholas Asher. 1993. *Reference to Abstract Objects in Discourse*, volume 50 of *SLAP*. Kluwer, <http://www.wkap.nl/>.
- Laura Banarescu, Claire Bonial, Shu Cai, Madalina Georgescu, Kira Griffitt, Ulf Hermjakob, Kevin Knight, Philipp Koehn, Martha Palmer, and Nathan Schneider. 2013. *Abstract Meaning Representation for sembanking*. In *Proceedings of the 7th Linguistic Annotation Workshop and Interoperability with Discourse*, pages 178–186, Sofia, Bulgaria. Association for Computational Linguistics.
- Adrien Barbaresi and Kay-Michael Würzner. 2014. For a fistful of blogs: Discovery and comparative benchmarking of republishable German content. In *Proceedings of the NLP4CMC Workshop (KONVENS 2014)*, pages 2–10. Hildesheim University Press.
- Douglas Biber. 1989. A typology of english texts. *Linguistics*, 27.1:3–43.
- Ann Bies, Mark Ferguson, Karen Katz, Robert MacIntyre, Victoria Tredinnick, Grace Kim, Mary Ann Marcinkiewicz, and Britta Schasberger. 1995. Bracketing guidelines for Treebank II style Penn Treebank project. Technical report, University of Pennsylvania.
- Steven Bird and Edward Loper. 2004. *NLTK: The natural language toolkit*. In *Proceedings of the ACL Interactive Poster and Demonstration Sessions*, pages 214–217, Barcelona, Spain. Association for Computational Linguistics.
- Sigmund von Birken. 1652. *Die Fried-erfreuete Teutonje*. Dümmler, Nürnberg.
- Abel Boyer. 1702. *The History of King William the Third: In III Parts*. Printed for A. Roper and F. Cogan, London.
- Chloé Braud, Yang Janet Liu, Eleni Metheniti, Philippe Muller, Laura Rivière, Attapol Rutherford, and Amir Zeldes. 2023. *The DISRPT 2023 shared task on elementary discourse unit segmentation, connective detection, and relation classification*. In *Proceedings of the 3rd Shared Task on Discourse Relation Parsing and Treebanking (DISRPT 2023)*, pages 1–21, Toronto, Canada. The Association for Computational Linguistics.
- Anne Breitbarth. 2005. *Live fast, die young: The short life of Early Modern German auxiliary ellipsis: Zugl.: Tilburg, Univ., Diss., 2005*, volume 115 of *LOT*. LOT, Utrecht.
- Lynn Carlson and Daniel Marcu. 2001. *Discourse tagging reference manual*. Technical Report ISI-TR-545, Information Sciences Institute (ISI), Marina del Rey, CA.
- Lynn Carlson, Daniel Marcu, and Mary Ellen Okurovsky. 2001. *Building a discourse-tagged corpus in the framework of Rhetorical Structure Theory*. In *Proceedings of the Second SIGdial Workshop on Discourse and Dialogue*.
- Claudia Claridge. 2025. *History writing*. In M. Kytö and E. Smitterberg, editors, *The New Cambridge History of the English Language: Documentation, Sources of Data and Modelling*, pages 433–458. Cambridge University Press, Cambridge.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. *Unsupervised cross-lingual representation learning at scale*. *Preprint*, arXiv:1911.02116.
- Marie-Catherine de Marneffe and Christopher D. Manning. 2008. *The Stanford typed dependencies representation*. In *Coling 2008: Proceedings of the workshop on Cross-Framework and Cross-Domain Parser Evaluation*, pages 1–8, Manchester, UK. Coling 2008 Organizing Committee.
- Marie-Catherine de Marneffe, Christopher D. Manning, Joakim Nivre, and Daniel Zeman. 2021. *Universal Dependencies*. *Computational Linguistics*, 47(2):255–308.
- David Denison. 1999. *SYNTAX*. In Suzanne Romaine, editor, *The Cambridge History of the English Language*, 1 edition, pages 92–329. Cambridge University Press.
- Explosion AI. 2025. *spacy: Industrial-strength natural language processing in python*.
- Steffen Frenzel, Maximilian Krupop, and Manfred Stede. 2026. *Discourse segmentation of german text*

- with pretrained language models. *Journal for Language Technology and Computational Linguistics*, 39(1):1–31.
- Annemarie Friedrich. 2017. *States, events, and generics: computational modeling of situation entity types*. Ph.D. thesis, Universität des Saarlandes.
- Annemarie Friedrich, Kleio-Isidora Mavridou, and Alexis Palmer. 2015a. *Situation Entity Types Annotation Manual*.
- Annemarie Friedrich and Alexis Palmer. 2014a. Automatic prediction of aspectual class of verbs in context. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 517–523, Baltimore, Maryland. Association for Computational Linguistics.
- Annemarie Friedrich and Alexis Palmer. 2014b. *Situation entity annotation*. In *Proceedings of LAW VIII - The 8th Linguistic Annotation Workshop*, pages 149–158, Dublin, Ireland. Association for Computational Linguistics and Dublin City University.
- Annemarie Friedrich, Alexis Palmer, Melissa Peate Sørensen, and Manfred Pinkal. 2015b. *Annotating genericity: a survey, a scheme, and a corpus*. In *Proceedings of the 9th Linguistic Annotation Workshop*, pages 21–30, Denver, Colorado, USA. Association for Computational Linguistics.
- Annemarie Friedrich, Alexis Palmer, and Manfred Pinkal. 2016. *Situation entity types: automatic classification of clause-level aspect*. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1757–1768, Berlin, Germany. Association for Computational Linguistics.
- Annemarie Friedrich and Manfred Pinkal. 2015a. Automatic recognition of habituais: a three-way classification of clausal aspect. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 2471–2481, Lisbon, Portugal. Association for Computational Linguistics.
- Annemarie Friedrich and Manfred Pinkal. 2015b. *Discourse-sensitive automatic identification of generic expressions*. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1272–1281, Beijing, China. Association for Computational Linguistics.
- Luke Gessler, Shabnam Behzad, Yang Janet Liu, Siyao Peng, Yilun Zhu, and Amir Zeldes. 2021. *DisCoDisCo at the DISRPT2021 shared task: A system for discourse segmentation, classification, and connective detection*. In *Proceedings of the 2nd Shared Task on Discourse Relation Parsing and Treebanking (DISRPT 2021)*, pages 51–62, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Jens Van Gysel, Meagan Vigus, Jayeol Chun, Kenneth Lai, Sarah Moeller, Jiarui Yao, Timothy J. O’Gorman, Andrew Cowell, W. Bruce Croft, Chu-Ren Huang, Jan Hajic, James H. Martin, Stephan Oepen, Martha Palmer, James Pustejovsky, Rosa Vallejos, and Ni-anwen Xue. 2021. *Designing a uniform meaning representation for natural language processing*. *KI - Künstliche Intelligenz*, 35:343 – 360.
- Nancy Ide, Collin Baker, Christiane Fellbaum, Charles Fillmore, and Rebecca Passonneau. 2008. *MASC: the manually annotated sub-corpus of American English*. In *Proceedings of the Sixth International Conference on Language Resources and Evaluation (LREC’08)*, Marrakech, Morocco. European Language Resources Association (ELRA).
- Justus-Liebig-Universität Gießen. 2022. Syntaktische grundstrukturen des neuhochdeutschen—korpus. <https://www.uni-giessen.de/de/fbz/fb05/germanistik/forschung/sprache/gieskane/korpus>. Accessed: 2026-03-04.
- Judith L. Klavans and Martin Chodorow. 1992. *Degrees of stativity: The lexical representation of verb aspect*. In *COLING 1992 Volume 4: The 14th International Conference on Computational Linguistics*.
- Ottolie Koralek. 1889/1890. *Lamentatio intermissa I*. Tagebucharchiv Emmendingen. Unpublished Transcription (Hollmann).
- Manfred Krifka, Francis Jeffrey Pelletier, Gregory N. Carlson, Alice ter Meulen, Godehard Link, and Genaro Chierchia. 1995. *Genericity: An introduction*. In Gregory N. Carlson and Francis Jeffry Pelletier, editors, *The Generic Book*, Studies in Communication, Media, and Public Opinion, pages 1–124. University of Chicago Press.
- Taku Kudo and John Richardson. 2018. *SentencePiece: A simple and language independent subword tokenizer and detokenizer for neural text processing*. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 66–71, Brussels, Belgium. Association for Computational Linguistics.
- John Lafferty, Andrew McCallum, and Fernando CN Pereira. 2001. *Conditional random fields: Probabilistic models for segmenting and labeling sequence data*. In *Proc. 18th International Conf. on Machine Learning.*, page 282–289.
- Sobha Lalitha Devi, Pattabhi Rk Rao, and Vijay Sundar Ram. 2025. *SeCoRel: Multilingual discourse analysis in DISRPT 2025*. In *Proceedings of the 4th Shared Task on Discourse Relation Parsing and Treebanking (DISRPT 2025)*, pages 79–86, Suzhou, China. Association for Computational Linguistics.
- Ilya Loshchilov and Frank Hutter. 2019. *Decoupled weight decay regularization*. *Preprint*, arXiv:1711.05101.

- David M. Magerman. 1995. [Statistical decision-tree models for parsing](#). In *33rd Annual Meeting of the Association for Computational Linguistics*, pages 276–283, Cambridge, Massachusetts, USA. Association for Computational Linguistics.
- William Mann and Sandra Thompson. 1988. Rhetorical structure theory: Towards a functional theory of text organization. *Text*, 8(3):243–281.
- William C. Mann, Christian M.I.M. Matthiessen, and Sandra A. Thompson. 1992. [Rhetorical Structure Theory and Text Analysis](#). In William C. Mann and Sandra A. Thompson, editors, *Pragmatics & Beyond New Series*, volume 16, page 39. John Benjamins Publishing Company.
- Thomas A. Mathew and Graham E. Katz. 2009. Supervised categorization for habitual versus episodic sentences. In *Proceedings of the Sixth Midwest Computational Linguistics Colloquium*, Bloomington, Indiana. Indiana University.
- Kleio-Isidora Mavridou, Annemarie Friedrich, Melissa Peate Sørensen, Alexis Palmer, and Manfred Pinkal. 2015. [Linking discourse modes and situation entity types in a cross-linguistic corpus study](#). In *Proceedings of the First Workshop on Linking Computational Models of Lexical, Sentential and Discourse-level Semantics*, pages 12–21, Lisbon, Portugal. Association for Computational Linguistics.
- Eleni Miltsakaki, Rashmi Prasad, Aravind Joshi, and Bonnie Webber. 2004. [The Penn Discourse Treebank](#). In *Proceedings of the Fourth International Conference on Language Resources and Evaluation (LREC'04)*, Lisbon, Portugal. European Language Resources Association (ELRA).
- Marc Moens and Mark Steedman. 1988. [Temporal ontology and temporal reference](#). *Computational Linguistics*, 14(2):15–28.
- Philippe Muller, Chloé Braud, and Mathieu Morey. 2019. [ToNy: Contextual embeddings for accurate multilingual discourse segmentation of full documents](#). In *Proceedings of the Workshop on Discourse Relation Parsing and Treebanking 2019*, pages 115–124, Minneapolis, MN. Association for Computational Linguistics.
- Kota Nayak. 2024. [Does chatgpt measure up to discourse unit segmentation? a comparative analysis utilizing zero-shot custom prompts](#). *Proceedings of the Canadian Conference on Artificial Intelligence*.
- Anna Nedoluzhko. 2013. [Generic noun phrases and annotation of coreference and bridging relations in the Prague dependency treebank](#). In *Proceedings of the 7th Linguistic Annotation Workshop and Interoperability with Discourse*, pages 103–111, Sofia, Bulgaria. Association for Computational Linguistics.
- Lev Pevzner and Marti A. Hearst. 2002. [A critique and improvement of an evaluation metric for text segmentation](#). *Computational Linguistics*, 28(1):19–36.
- Livia Polanyi. 1995. The linguistic structure of discourse. Csl technical report, CSLI, Stanford, CA.
- Livia Polanyi, Chris Culy, Martin van den Berg, Gian Lorenzo Thione, and David Ahn. 2004. [A rule based approach to discourse parsing](#). In *Proceedings of the 5th SIGdial Workshop on Discourse and Dialogue at HLT-NAACL 2004*, pages 108–117, Cambridge, Massachusetts, USA. Association for Computational Linguistics.
- Thomas Salmon. 1724. *A Review of the History of England. In Two Volumes*, 2 edition, volume 1. Printed for Charles Rivington, London.
- Carlota S Smith. 2003. *Modes of discourse: The local structure of texts*, volume 103. Cambridge University Press.
- Radu Soricut and Daniel Marcu. 2003. [Sentence level discourse parsing using syntactic and lexical information](#). In *Proceedings of the 2003 Human Language Technology Conference of the North American Chapter of the Association for Computational Linguistics*, pages 228–235.
- William Stubbs. 1891. *The Constitutional History of England, Vol. I*. Oxford at the Clarendon Press, Oxford. Digitized by the Digital Library of India (Osmania University).
- Victoria Thomas. 2018. *Auxiliary Ellipsis in Early Modern German 1350-1800*. Phd, The University of Manchester, Manchester.
- Erik F. Tjong, Kim Sang, and Hervé Déjean. 2001. Introduction to the CoNLL-2001 shared task: Clause identification. In *Proceedings of the 2001 Workshop on Computational Natural Language Learning (CoNLL)*, Manchester, UK.
- Milan Tofiloski, Julian Brooke, and Maite Taboada. 2009. A syntactic and lexical-based discourse segmenter. In *Proceedings of the 47th Annual Meeting of the Association for Computational Linguistics and the 4th International Joint Conference on Natural Language Processing of the AFNLP (ACL-IJCNLP)*, pages 77–80, Singapore.
- Nikki Van De Pol and Peter Petré. 2015. [Why is there a Present-Day English absolute?](#) *Studies in Language*, 39(1):199–229.
- Zeno Vendler. 1957. *Linguistics in Philosophy*, chapter Verbs and Times. Cornell University Press, Ithaca, New York.
- Yizhong Wang, Sujian Li, and Jingfeng Yang. 2018. [Toward fast and accurate neural discourse segmentation](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 962–967, Brussels, Belgium. Association for Computational Linguistics.

Bonnie Webber, Rashmi Prasad, Alan Lee, and Aravind Joshi. 2019. The penn discourse treebank 3.0 annotation manual. *Philadelphia, University of Pennsylvania*, 35:108.

Egon Werlich. 1989. *Typologie der Texte*. UTB für Wissenschaft.

## A Discourse Segmentation Guidelines in Related Work

This section takes a closer look at related work on discourse segmentation including more extensive examples. In the Linguistic Discourse Model (LDM; Polanyi (1995)), the atomic units of discourse are called Basic Discourse Units (BDUs); they are segments that have the potential to establish an anchor point for future attachment of other segments. Polanyi et al. (2004) start by defining the semantic basis for functioning as a segment and then identify syntactic constructions that are able to carry the semantic information needed for discourse segment status. They observe that in written text, often a subsequent but not necessarily adjacent segment continues the development of material introduced in a sub-sentential, often subordinate, clause. Example (23) shows an LDM annotation which treats the post-modifier “braying” next door as a separate segment, as it can be interpreted as “who was braying next door.”

(23) [The donkey [braying next door] was loud.]

Nominal gerunds “Singing” and nominalizations “rationalization” are not always considered BDU segments. However, Polanyi et al. consider the nominal “the destruction of the old town hall” to be a BDU. To the best of our knowledge, no guidelines exist for distinguishing these cases.

Rhetorical Structure Theory (RST; Mann et al. (1992)), presents another approach to discourse segmentation. Here, the segments are called Elementary Discourse Units (EDUs). EDUs are defined essentially as clauses, but clausal subjects, complements and restrictive relative clauses are considered as parts of the clause headed by their governing verb. When building the RST Discourse Treebank, Carlson et al. (2001) note that applying this intuitive notion is difficult when aiming for a large and consistently annotated corpus. They develop an extensive set of rules for identifying EDUs based on syntactic constituents with the aim of obtaining a balance between tagging granularity and the ability to identify units consistently (Carlson and Marcu, 2001). The rules are motivated by RST’s

inventory of discourse relations (schemas). For example, while infinitival constructions are generally not considered to constitute their own EDU, they are if they introduce a purpose clause as in (24) because the infinitival clause corresponds to the satellite of a Purpose relation here. Prepositional phrases with clausal objects (25) are EDUs, while other non-finite clausal objects are not.

(24) [... officials at Southern Co. conspired to cover up their accounting for spare parts] [to evade federal income taxes.]

(25) [Canadian Utilities isn’t alone] [in exploring power generation opportunities in Britain.]

For ease of comparison, we provide an overview of the differences between SE segmentation and other discourse segmentation schemes in Figure 2. A core difference is that SE segmentation is motivated by the ability of assigning aspectual features, i.e., situation entity types (Smith, 2003), to a segment, while RST, LDM, and PDTB are more concerned with whether and what discourse relations can be identified between segments. Segmentation in the PDTB can be either finer-grained or coarser-grained than SE segmentation; LDM and RST apply either the same or a more fine-grained segmentation as illustrated in Figure 2. Connectives typically constitute a separate data type and are not included in the arguments, while SE segmentation performs an exhaustive text segmentation, simply grouping them with the clause they introduce.

## B Background on Linguistic Framework and Annotation Guidelines

Table 5 shows the full inventory of SE types including their description. SE annotation into these classes is a downstream task building on SE segmentation. In a similar vein, Table 6 defines the relationship between the main verb’s and the main referent’s aspectual features and the SE type of the entire segment.

## C Manually Annotated Data

Table 7 shows the data source for each manually annotated file as well as its publication year. Token counts are available in Table 2. Figure 3 shows the inter-annotator disagreement analysis across annotation conditions and linguistic categories. Darker cells indicate higher disagreement frequencies.

SE type	Description	Example
<b>Eventualities</b>		
STATE	introduce properties	The colonel owns the farm.
EVENT	happenings	John won the race.
REPORT	for attribution	"...", said Obama.
<b>General Statives</b>		
GENERIC SENTENCE	generalizations over kinds	The lion has a bushy tail.
GENERALIZING SENTENCE	habituals; generalizations over situations	Mary often fed the cat last year.
<b>Abstract Entities</b>		
FACT	clausal complements of verbs of knowledge	I know that she refused the offer.
Proposition	clausal complements of verbs of belief	I believe that she refused the offer.
QUESTION		Who wants to come?
IMPERATIVE		Hand me the pen!

Table 5: Inventory of SE types, as adapted from Smith (2003) in previous work (Friedrich et al., 2016; Friedrich and Palmer, 2014b).

SE type	main referent	aspectual class	habituality
EVENT	non-generic	dynamic	episodic
	generic		
STATE	non-generic	stative	static
GENERIC SENTENCE	generic	dynamic	habitual
		stative	static, habitual
GENERALIZING SENTENCE	non-generic	dynamic	habitual
		stative	

Table 6: SE types and their corresponding features, adapted from Friedrich (2017).

Variety	Text	Source	Year
Contemp. E*	Email	MASC	1993
	News WSJ0135	MASC	>1990
	Travelblog Dublin	MASC	<2008
	Trees	Wikipedia	<2016
LModE	Historiography - Boyer	CLMEH	1702
	Historiography - Salmon	CLMEH	1736
Contemp. G	bitblokes	Blog	2023
	phantanews	Blog	2026
	iphoneblog	Blog	2025
	literaturblog	Blog	2017
	Literatur	Wikipedia	2026
	Wald	Wikipedia	2026
NHG	Lecture - Anthus	GiesKaNe	1838
	Academic - Becher	GiesKaNe	1668

Table 7: Text sources used in the annotation study. \*Contemporary E texts are excluded from model training.

## D Computational Modeling

Training is performed using the AdamW optimizer (Loshchilov and Hutter, 2019) with default hyperparameters  $\beta_1 = 0.9$ ,  $\beta_2 = 0.999$ ,  $\epsilon = 1 \times 10^{-8}$ , a batch size of 64, a maximum of 20 epochs, and early stopping after three non-improving epochs based on the boundary (B) tag’s F1 score. A small

grid search is conducted over five random seeds, learning rates ( $4e-5$ ,  $5e-5$ ), and weight decay values (0.001, 0.005) to select the best model. We use linear learning rate decay with a 10% warmup phase, gradient clipping with a maximum norm of 1.0, and BF16 mixed precision training on four NVIDIA H100 GPUs. The best individual run reached an F1 of 90.65, converging after 8 epochs with a learning rate of  $4 \times 10^{-5}$  and weight decay 0.005, random seed 100.

**WindowDiff** We calculate the WindowDiff (Pevzner and Hearst, 2002) using the NLTK implementation (Bird and Loper, 2004) as follows: Given a sentence of length  $N$  with reference boundaries  $r_i$  and predicted boundaries  $h_i$  (where  $r_i = 1$  if token  $i$  starts a new segment, 0 otherwise), we:

1. Set the window size to

$$k = \max\left(1, \left\lfloor \frac{N}{2B} \right\rfloor\right),$$

where  $B$  is the total number of reference boundaries in the sentence.

2. Slide a window of size  $k$  across the sentence, comparing the boundary count in each win-

PDTB vs. SE:	... [ and many are hoping [ for major new liberalizations ]ARG2]SE [ if [ he is returned firmly to power. ]ARG1]SE
PDTB vs. SE:	[ [ Here in this new center for Japanese assembly plants just across the border from San Diego,]SE [ turnover is dizzying, infrastructure shoddy, bureaucracy intense. ]SE [ Even after-hours drag;]SE [ “karaoke” bars, ]SE† [ where Japanese revelers sing over recorded music,]SE [ are prohibited by Mexico’s powerful musicians union. ]SE ]ARG2 [ Still, [ 20 Japanese companies, ]SE† [ including giants such as Sanyo Industries Corp., Matsushita Electronics Components Corp. and Sony Corp. ]SE [ have set up shop in the state of Northern Baja California. ]SE ]ARG1
LDM vs. SE:	[ [ California elected Schwarzenegger ]BDU [ governor ]BDU. ] SE
LDM vs. SE:	[ [ The donkey ]SE† [ [ braying next door ]SE ] [ was annoying. ]SE ]BDU
RST vs. SE:	[ [ The company announced ]SE]EDU [ [ that it will shut down its plant ]SE]EDU [ [ and dismiss several hundred employees. ]SE]EDU
RST vs. SE:	[ [ The company plans to shut down its plant ]SE [ and dismiss several hundred employees.]SE]EDU

Figure 2: Comparison of SE segmentation to RST, LDM, and PDTB. † are examples of segments that are not complete SEs on their own, but they also do not belong to the following situation (unless by coreference relations). As illustrated in Figure 1, they can still contain main referents of other segments and hence be part of another segment as defined by syntactic dependency structure.

dow:

$$WD = \frac{1}{N-k} \sum_{i=1}^{N-k} \mathbf{1} \left( \sum_{j=i}^{i+k-1} r_j \neq \sum_{j=i}^{i+k-1} h_j \right)$$

The score ranges from 0 (perfect match) to 1 (complete mismatch). Unlike exact match, WindowDiff tolerates small boundary shifts: if a boundary is off by one or two positions, only the windows containing that region are penalized, not the entire sentence.

We compute WindowDiff per sentence and average across all sentences with  $N \geq 2k + 1$ , where  $M$  is the number of such sentences:

$$WD_{\text{final}} = \frac{1}{M} \sum_{s=1}^M WD^{(s)}.$$

## E Additional Annotation Peculiarities and Annotation Examples

### E.1 Late Modern English

**Relative clauses** Segmentation decisions are affected by variation in relative clause formation that can be observed in Late Modern English. The use of *which* for human antecedents is declining in this period and zero relatives are used more commonly in their place (Denison, 1999), which creates ambiguity in identifying clause boundaries. Examples (26a) and (26b), taken from Stubbs (1891),

illustrate this variation. Example (26b) shows the special case of reduced relative clauses in the past tense as discussed in the general guidelines for English.

- (26) (a) [The Saxons, Angles, and Jutes,][ although speaking the same language,][ worshipping the same gods][ and using the same laws,][ had no political unity like the Franks of Clovis;]  
(b) [the Saxons in Germany were still a pure nationality,][unconquered by the Franks,][untainted by Roman manners,][ and still heathen.]

Since the Late Modern English period covers an ongoing usage change with regards to relative clauses and their positions, the guideline for using preposition or post-position to approximate adjectival use cannot be relied upon. Segmentation, instead, is based on the annotators perception given the larger semantic embedding.

### E.2 Late Modern English

Figure 4 reveals how the model treats an entire sentence as a single segment, failing to introduce a boundary at the transition to the main clause. This error highlights a limitation of the model in handling subordination structures, particularly in historical German, where clause boundaries may be less clearly signaled by punctuation or conjunctions. Whilst humans reliably segment along clause

		orthography				syntax			subordination				non finite			style				discourse		other		
<b>DE_CURRENT</b>	0	6	3	3	0	25	5	1	14	0	4	5	0	0	0	0	14	9	0	0	0	0	0	12
<b>DE_HIST</b>	1	3	0	0	2	50	0	6	34	0	9	4	0	0	0	11	9	0	11	0	0	2	18	
<b>E_CURRENT</b>	0	0	0	0	0	30	4	20	12	0	0	0	11	6	21	2	9	0	2	0	0	2	34	
<b>E_HIST</b>	0	0	0	0	0	30	0	18	14	2	0	4	0	15	0	4	6	0	14	32	2	0	9	
		comma	dash	slash	dotdotdot	quotes	clause missed	sentence boundary	verbless	relative clause	reduced relative clause	conditional	comparison	gerund	to infinitive	participle	list	inlay	embedded example	ellipsis	absolute construction	modality	reported speech	arbitrary boundary

Figure 3: Inter-annotator disagreement analysis across annotation conditions and linguistic categories. Darker cells indicate higher disagreement frequencies.

boundaries, the model under-segments here. Figure 5 shows one further NHG example that illustrates a genuinely non-trivial case due to ellipses and multiple coordinated and subordinated structures where inter annotator differences can be seen as valid alternative annotations.

### E.3 Contemporary German

Figure 6 illustrates a case of annotator segmentation disagreement involving the German coordinating conjunction “und.” Annotator 1 and the model introduce a boundary at the conjunction, treating the coordinated clauses as separate SEs, whereas Annotator 2 merges the coordinated clauses into a single segment. The example highlights a linguistically motivated source of variability in SE annotation: coordinated clauses can plausibly be interpreted either as independent segments or as a unified discourse unit.

Token <i>English</i>	Daß <i>That</i>	bei <i>in</i>	den <i>the</i>	Pflanzen <i>plants</i>	von <i>of</i>	keinem <i>no</i>	eigentlichen Essen <i>actual food</i>
Model	B	I	I	I	I	I	I
Annotator 1	B	I	I	I	I	I	I
Annotator 2	B	I	I	I	I	I	I

Token <i>English</i>	die <i>the</i>	Rede <i>talk</i>	sein <i>be</i>	kann <i>can</i>	,	<b>wird</b> <b>will</b>	sich <i>itself</i>	aus <i>from</i>
Model	I	I	I	I	I	<b>I</b>	I	I
Annotator 1	I	I	I	I	I	<b>B</b>	I	I
Annotator 2	I	I	I	I	I	<b>B</b>	I	I

Token <i>English</i>	dem <i>the</i>	Begriffe <i>concept</i>	des <i>of</i>	Essens <i>eating</i>	später <i>later</i>	ergeben <i>result</i>	.
Model	I	I	I	I	I	I	I
Annotator 1	I	I	I	I	I	I	I
Annotator 2	I	I	I	I	I	I	I

Figure 4: Example of segmentation disagreement on a subordinated clause. The model treats the entire sentence as a single segment, while human annotators introduce a segment boundary at the main clause (“wird”).

Token <i>English</i>	[...] <i>,</i>	sehen <i>we</i>	wir <i>see</i>	auf <i>on</i>	jener <i>that</i>	ersten <i>first</i>	Stufe <i>level</i>	Wesen <i>beings</i>	,	welche <i>which</i>	,	mit <i>with</i>
Model	I	I	I	I	I	I	I	I	I	<b>B</b>	I	I
Annotator 1	I	I	I	I	I	I	I	I	I	<b>B</b>	I	I
Annotator 2	I	I	I	I	I	I	I	I	I	<b>B</b>	I	I

Token <i>English</i>	Ausnahme <i>exception</i>	des <i>of</i>	Salzes <i>salt</i>	und <i>and</i>	Wassers <i>water</i>	,	weder <i>neither</i>	eßbar <i>edible</i>	sind <i>are</i>	noch <i>nor</i>	essen <i>eat</i>	;
Model	I	I	I	I	I	I	I	I	I	I	I	I
Annotator 1	I	I	I	I	I	I	I	I	I	I	I	I
Annotator 2	I	I	I	I	I	I	I	I	I	I	I	I

Token <i>English</i>	auf <i>on</i>	der <i>the</i>	zweiten <i>second</i>	eßbare <i>edible</i>	,	aber <i>but</i>	nicht <i>not</i>	essende <i>eating</i>	Wesen <i>beings</i>	;	auf <i>on</i>	der <i>the</i>
Model	I	I	I	I	I	I	I	I	I	I	I	I
Annotator 1	<b>B</b>	I	I	I	I	I	I	I	I	I	I	I
Annotator 2	I	I	I	I	I	I	I	I	I	I	I	I

Token <i>English</i>	dritten <i>third</i>	Stufe <i>level</i>	endlich <i>finally</i>	Wesen <i>beings</i>	,	welche <i>which</i>	essen <i>eat</i>	und <i>and</i>	gegessen <i>are eaten</i>	werden <i>become</i>	.
Model	I	I	I	I	I	I	I	I	I	I	I
Annotator 1	I	I	I	I	I	<b>B</b>	I	<b>B</b>	I	I	I
Annotator 2	I	I	I	I	I	I	I	I	I	I	I

Figure 5: Example of segmentation disagreement in a complex sentence with multiple coordinated and subordinated structures.

Token	Die	App	unterstützt	natürliche	Spracheingabe
<i>English</i>	<i>The</i>	<i>App</i>	<i>supports</i>	<i>natural</i>	<i>speech recognition</i>
Model	B	I	I	I	I
Annotator 1	B	I	I	I	I
Annotator 2	B	I	I	I	I
Token	<b>und</b>	bietet	intelligente	Vorschläge	.
<i>English</i>	<b><i>and</i></b>	<i>offers</i>	<i>intelligent</i>	<i>suggestions</i>	.
Model	<b>B</b>	I	I	I	I
Annotator 1	<b>B</b>	I	I	I	I
Annotator 2	I	I	I	I	I

Figure 6: Example of segmentation disagreement on a coordinated clause. Annotator 1 and the model introduce a new segment boundary at “und”, while Annotator 2 treats the entire sentence as a single segment.

# UD-CHILDES-BG: a dependency treebank of Bulgarian child and child-directed speech

Mila Marcheva-Nash<sup>1</sup> Yasena Chantova<sup>2,3</sup> Tsvetina Kirilova<sup>2</sup> Ivelina Pavlova<sup>2</sup>  
Tsvetelina Stefanova<sup>2</sup> Yoana Vasileva<sup>2</sup> Weiwei Sun<sup>1</sup>

<sup>1</sup>Department of Computer Science & Technology, University of Cambridge

<sup>2</sup>Faculty of Slavic Studies, Sofia University “St. Kliment Ohridski”

<sup>3</sup>University of Library Studies and Information Technologies, Bulgaria

## Abstract

This paper presents (i) UD-CHILDES-BG, a manually corrected Universal Dependencies treebank of Bulgarian child and child-directed speech, (ii) a quantitative and phenomenon-based evaluation of inter-annotator agreement on developmental data, and (iii) a systematic analysis of parser errors in this underrepresented domain. We manually correct 4,338 dependency parses (10% of the CHILDES-BG corpus), of which 14% are double-annotated. Inter-annotator agreement on UAS/LAS is 91.71/86.12 for child-directed speech (CDS) and 88.14/81.40 for child speech (CS). Parser performance on the manually corrected portion is 92.70/85.54 for CDS and 90.97/81.52 for CS, compared to a reported 93.37/90.21 on the test set of adult written language. Our analyses reveal that CDS and CS pose challenges for dependency annotation and parsing, particularly in discourse-related structures, which are less common in adult written language.

## 1 Introduction

Linguistically annotated corpora are central to supervised NLP, evaluation of LLMs, and theory-building in linguistics. Most syntactic treebanks are based on adult written language, while developmental data remains comparatively underrepresented, despite growing interest in the role of syntactic structure for language model robustness (Güven et al., 2025). Furthermore, first language acquisition (FLA) research relies on child-directed speech (CDS) and child speech (CS) with high quality annotations to make extrapolations about the development of language (Bosch, 2025). Providing linguistically annotated CS and CDS, along with scalable guidelines, for resource-poor languages is thus a useful area of expansion to further FLA and NLP research.

The challenge of parsing CDS and/or CS was first addressed by Sagae et al. (2001), and has

since been explored primarily for English (Pearl and Sprouse, 2013; Liu and Prud’hommeaux, 2023; Yang et al., 2025), with only a small number of studies on other languages, including Hebrew (Szubert et al., 2024), Japanese (Butler et al., 2022), and Dutch (Odiijk et al., 2018). Universal Dependencies (UD; de Marneffe et al., 2021) provides a cross-linguistically consistent framework for syntactic annotation, which has been applied to over one hundred languages. While treebanks for CS and CDS are scarce, they are mainly within the UD framework. For Slavic and Balkan languages, no manually corrected treebank for developmental data currently exists.

In this paper, we present UD-CHILDES-BG, a treebank of manually corrected UD parses for Bulgarian. UD-CHILDES-BG and the accompanying analysis scripts are available on GitHub.<sup>1</sup> The annotation scheme is supplied in Appendix B. We manually correct 4,338 dependency parses, corresponding to 10% of the CHILDES-BG corpus (Popova and Popov, 2020), with a breakdown of 2,481 CDS and 1,857 CS utterances. Starting from automatically produced UD parses, we perform targeted manual correction of core layers: lemma, UPOS, morphological features, dependency arc, and dependency relation.

We assess annotation reliability through double annotation of 14% of the data: inter-annotator agreement on UAS/LAS is 91.71/86.12 for CDS and 88.14/81.40 for CS. By treating annotation disagreement as linguistically informative, we provide annotation guidance and linguistic background for a selected set of language phenomena, including clitic doubling, vocative case, and reflexive pronoun, which are prominent in CDS and CS, but less frequent in adult written language.

We further evaluate the parser on the manually

<sup>1</sup><https://github.com/milamarcheva/UD-CHILDES-BG>

corrected parses: UAS/LAS is 92.70/85.54 for CDS and 90.97/81.52 for CS, compared to a reported<sup>2</sup> 93.37/90.21 on a test set of adult written language. A per-relation analysis shows that the decreased parser performance on CS and CDS is due to discourse-specific structures frequent in conversational and developmental data, but less common in adult written language on which the parser has been trained.

## 2 Background

Bulgarian is a South Slavic language that exhibits both core Slavic properties, such as diminutives in CDS (Kempe and Brooks, 2001), and characteristic Balkan Sprachbund features, including an atrophied nominal case system and clitic doubling (Tomić, 2011). Bulgarian is *pro*-drop and displays relatively flexible word order, with agreement morphology on the verb marking person and number. These properties interact with discourse structure and clitic placement, making attachment decisions surrounding the phenomena less straightforward in dependency annotation. In CS and CDS such phenomena are further amplified by fragmentary utterances, vocative forms, and discourse particles.

Several syntactic resources exist for Bulgarian. BulTreeBank (BTB; Simov et al., 2002b,a; Simov and Osenova, 2003; Simov et al., 2004; Osenova and Simov, 2004; Simov et al., 2005) is a constituency treebank based on head-driven phrase structure grammar (HPSG; Pollard and Sag, 1994). UD-BTB (Osenova and Simov, 2017) is the manually corrected dependency resource based on BTB. The Bulgarian CHILDES corpus (CHILDES-BG) provides longitudinal developmental data but does not come with syntactic annotation (Popova and Popov, 2020). Parsing tools for Bulgarian include a Berkeley constituency parser trained on BTB (Petrov et al., 2006), as well as dependency parsers such as Stanza (Qi et al., 2020) and CLASSLA-Stanza (Terčon and Ljubešić, 2023) trained on UD-BTB. With regards to resources focusing on child Bulgarian, the existing literature covers a version of the MacArthur–Bates Communicative Development Inventory for Bulgarian (Andonova, 2015), as well as studies on the development of grammar in Bulgarian (Popova and Filipov, 2022; Popova, 2023).

UD resources for CS and CDS remain limited

<sup>2</sup><https://stanfordnlp.github.io/stanza/performance.html>

cross-linguistically, starting with the foundational dependency treebank for only CDS by Sagae et al. (2001, 2010). More recently, dependency treebanks covering both CS and CDS have been developed for English (Liu and Prud’hommeaux, 2023; Yang et al., 2025), Japanese (Butler et al., 2022), and Hebrew (Gretz et al., 2013; Szubert et al., 2024), with some cross-linguistic work on English–Hebrew corpora (Szubert et al., 2024). Constituency resources are even scarcer, including the English CHILDES-TB (Pearl and Sprouse, 2013) only for CDS and the Dutch AnnCor Treebank (Odijk et al., 2018) for both CS and CDS. Finally, there are efforts for providing automatic UD annotation to all available CHILDES transcripts (MacWhinney, 2012; Liu, 2024) via `batchalign`.<sup>3</sup> However, at the time of release of `batchalign`, the Bulgarian CHILDES corpus was published only with Latin transcription, and the UD parsing tools require Cyrillic input. The authors of the Bulgarian CHILDES corpus have since uploaded a Cyrillic transcription, and we independently transliterated the utterances for this project. Still, no manually corrected treebank currently exists for Bulgarian CS or CDS, or any other Balkan or Slavic language.

We adopt the UD framework for three reasons. First, UD provides cross-linguistically consistent guidelines that enable direct comparison with existing CS and CDS resources. Second, the primary constituency resource for Bulgarian, BTB, employs relatively flat phrase-structure representations (Osenova and Simov, 2004), making UD’s dependency representation comparable in structural granularity while avoiding commitments to language-specific phrase-structural conventions. Third, UD annotation is directly compatible with widely used parsing toolkits facilitating error analysis and automatic pre-annotation. In choosing UD, we prioritise reproducibility and alignment with existing developmental treebanks.

Beyond resource creation, recent work in annotation research has emphasized the importance of distinguishing between annotator disagreement and objectively incorrect annotation (Klie et al., 2023; Weber-Genzel et al., 2024). However, systematic analyses of annotation error in developmental and non-canonical spoken data remain rare. Our work contributes to this line of research by examining parser error and annotation contention specifically in Bulgarian CS and CDS, a typolo-

<sup>3</sup><https://github.com/TalkBank/batchalign2>

gically and discourse-rich setting that challenges standard UD assumptions.

### 3 Data and procedure

#### 3.1 Source Data

We use the longitudinal section of the CHILDES-BG corpus (Popova and Popov, 2020), which follows five children between the ages of one and three years and covers both the child-directed speech and the first productions of the target children. In total there are 45,000 utterances, but after removing utterances with empty transcription, 43,915 remain, see Table 1 for a breakdown by child.

Child	CS		CDS		%
	M	T	M	T	
ALE	379	3682	430	4144	10.3
BOG	97	961	148	1465	10.1
ELI	121	1208	706	6825	10.3
SIM	924	9540	805	8347	9.7
TEF	336	3586	392	4157	9.4
Total	1857	18977	2481	24938	9.9

Table 1: Breakdown of valid (non-empty) sentences in the CHILDES-BG corpus by target child. Columns show the number of manually corrected sentences (M) and the total number of valid sentences (T) for child speech (CS) and child-directed speech (CDS). % denotes what fraction of total (CS+CDS) sentences for a given child were manually corrected.

A stratified sampling strategy was adopted across children and speech types, CS and CDS. Sampling was stratified by child and by age, and sentences were drawn from across the full set of available utterances, in order to ensure broad coverage of speakers and developmental stages. The primary annotator corrected slices across all children and both registers, while five additional annotators were each assigned specific non-overlapping subsets (see Appendix A for further detail). To ensure quality control, portions of the data were deliberately double-annotated. The exact numbers of sentences annotated for each target child are presented in Table 1. In total, 1,857 CS utterances and 2,481 CDS utterances were manually corrected, corresponding to 10% of all valid CS and CDS utterances in the corpus. Double annotation was performed on 14% of these.

#### 3.2 Preprocessing of CHILDES annotation

CHILDES corpora are transcribed and annotated following the CHAT guidelines<sup>4</sup> (MacWhinney, 1992). The CHAT guidelines cover child-language-specific cases such as phonological variation, e.g. *popo* [: *hippopotamus*], where the bracketed expression is the standard form of the child form preceding it, or special coding for unintelligible words, which is standardly *xxx*. Sentences consisting entirely of *xxx* are dropped and do not appear in the final dataset. However, when *xxx* is part of a sentence, we retain such sentences and annotate them following Odijk et al. (2018). Because we are using automatic parses for the initial preprocessing of the utterances, we perform normalisation where phonological variation annotation is provided: we replace the child form with the adult form in order to allow for better automatic parsing.

The CHILDES annotation is not always consistently applied, which poses a challenge to automatic cleaning of the annotation. The main focus of the manual data normalisation is the phonological variation annotation. Sometimes it is applied consistently as in Figure 1 (a), where automatic normalisation is sufficient.

Дугата [: другата] ана [: страна]  
dugata [: drugata] ana [: strana]  
other.DEF side  
‘the other side’  
**Normalised:** drugata strana

(a) Consistent use of CHAT annotation.

икам [: искам] гая [: да играя]  
ikam [: iskam] gaja [: da igraja]  
want.1SG play.1SG  
‘I want to play’  
**Problem:** inserted *da*

(b) Inconsistent use of CHAT annotation.

Figure 1: Use of CHILDES correction brackets: (a) consistent phonological normalisation; (b) inconsistent annotation introducing additional syntactic material

However, in other cases the bracketed notation is only applied to correct one of several repeated forms, or maps one surface form to several adult forms, as is the case in Figure 1 (b), where the notation is used inappropriately to introduce the functional token *da*. Although *da* is implied in the meaning, it should not be added to the child utterance as this utterance is representative of a specific

<sup>4</sup><https://talkbank.org/info/manuals/CHAT.html>

stage in FLA development, where function words are omitted. During the manual normalisation of the data we ensure that child utterances are corrected where possible with the word forms to facilitate automatic parsing, but we do not allow for extra syntactic material to be added.

### 3.3 Annotation procedure

We use a custom-made project on the INCEpTION platform (Klie et al., 2018). INCEpTION has built-in support for data in CoNLL format, which is the standard UD format. The annotation correction includes: correction of lemma, UPOS, morphological features, dependency arcs, and dependency relations. The UD-BTB also has a very fine-grained POS tag, XPOS, with encoded morphological information (Simov et al., 2004). For UD-CHILDES-BG, the XPOS tag is not corrected or retained, as the UPOS and morphological features overlap with it. The valid (non-empty) sentences from Table 1 were all parsed using the Stanza parser (Qi et al., 2020) trained on UD-BTB. The parses of the sentences selected via stratified sampling were uploaded to INCEpTION for annotation.

The data is annotated by six annotators, who are all native speakers of Bulgarian. A1 is the main annotator who created the annotation schemes, and annotated a total of 2,500 of the 4338 annotated sentences. The annotation agreement scores are calculated between A1 and the other annotators. A1 has a background in computational linguistics including a completed undergraduate course on *Formal Models of Language* and graduate courses *Natural Language Processing (NLP)*, *Introduction to Computational Semantics*, *Introduction to Natural Language Syntax and Parsing*, as well as teaching *Formal Models of Language* to undergraduate students. A2-A6 all have bachelor degrees in various philologies, and are currently enrolled in a Computational Linguistics Master’s programme. All annotators were previously familiar with the UD framework, and received a 2-hour training and overview of UD examples specific for this project, additional to previous experience.

The annotators first annotated 5% of the sentences, 1,204 CDS and 907 CS, with 238 CDS and 185 CS double annotated (all by A1, and non-overlapping subsets by the other annotators). Afterwards the annotations were analysed to identify the most common sources of error in annotation, and a 2-hour discussion was held among the annotators to decide on a unified approach for the edge cases.

Following the conclusions of the discussion, the annotation scheme was refined, and the annotators revised their initial annotations and annotated a further 5% of the data following the unified principles. The precise breakdown of sentences by annotator can be found in Appendix A.

## 4 Annotator agreement

Below we present the annotator agreement with standard quantitative metrics and further discuss specific phenomena worth mentioning due to their pertinence in CS and CDS. Sentences that deviate from written adult language, due to exhibiting features of spoken or developmental language, are not fully captured by the existing gold standard, UD-BTB. We provide further detail on some prominent examples below, while a more comprehensive list of challenging phenomena is covered in Appendix B.

### 4.1 Quantitative evaluation

We use several metrics to quantify inter-annotator agreement: UPOS and lemma accuracy measure exact agreement between annotators on universal part-of-speech tags and lemma assignment respectively; Cohen’s  $\kappa$  measures agreement beyond chance for UPOS annotation; Unlabelled Attachment Score (UAS) measures agreement on dependency head attachment regardless of relation label; and Labelled Attachment Score (LAS) measures agreement on both dependency head attachment and dependency relation labels. In the context of inter-annotator agreement, “accuracy” refers to pairwise agreement between annotators and does not imply comparison against an external gold-standard annotation.

Table 2 displays inter-annotator agreement on the double-annotated subsets of CDS and CS. Overall agreement is high across domains. For CDS, UPOS accuracy reaches 92.48 ( $\kappa = .9143$ ), while CS yields slightly lower but comparable values 90.57 ( $\kappa = .8899$ ). Lemma agreement is lower in both domains (86.50 for CDS; 79.25 for CS), reflecting increased ambiguity in lemmatisation of non-canonical and morphologically reduced forms. Syntactic attachment agreement remains strong in both domains, with UAS/LAS of 91.71/86.12 for CDS and 88.14/81.40 for CS. The similarity of LAS across CDS and CS suggests that, despite structural irregularities in child speech, annotators converge on dependency labels at comparable rates

once head attachment is established. An additional explanation of the convergence of UAS/LAS scores for CS and CDS is that CS is naturally comprised of shorter utterances, hence, there are fewer arcs and relations the annotators can disagree on.

	$UPOS_a$	$UPOS_\kappa$	Lemma	UAS	LAS
CDS	92.48	0.9143	86.50	91.71	86.12
CS	90.57	0.8899	79.25	88.14	81.40

Table 2: Combined inter-annotator agreement: UPOS (accuracy and Cohen’s  $\kappa$ ); Lemma (accuracy); Unlabelled Attachment Score (UAS) and Labelled Attachment Score (LAS).

Table 3 breaks down agreement by dependency relation.  $IAA$  (inter-annotator agreement) is calculated using the F1 formula. Structural relations such as *case*, *root*, and *aux* exhibit consistently high  $IAA$  in both domains, indicating stable annotation of core grammatical structure. In contrast, discourse-sensitive relations show greater variability. In CDS, *discourse* ( $IAA=84.82$ ) and *iobj* ( $IAA=80.43$ ) show reduced agreement, reflecting ambiguity in clitic doubling and particle attachment. In CS, agreement for *discourse* drops substantially ( $IAA=55.56$ ), alongside lower scores for *conj* and *obj*, suggesting that developmental constructions and fragmentary utterances increase annotation difficulty.

These patterns indicate that disagreement is concentrated not in core syntactic relations but in constructions that interact with discourse structure, clitic systems, and child-specific production phenomena. Rather than treating such disagreement as annotation noise, we interpret it as evidence of structural tension between standard UD guidelines and non-canonical spoken Bulgarian.

## 4.2 Phenomenon-based evaluation

Below we discuss annotators’ interpretations of several phenomena illustrative of CDS and CS in Bulgarian, using linguistic background to justify the final annotation decision. For all cases we follow the logic of annotation of UD-BTB (Osenova and Simov, 2017). In the interest of space we limit the discussion to the annotation of dependency arcs and relations; see the annotation scheme in Appendix B for comments on lemma, UPOS, and morphological features.

CDS		CS	
Relation	$IAA$	Relation	$IAA$
root	93.66	root	95.88
obj	91.03	obj	70.83
advmod	93.77	conj	62.22
nsubj	86.19	advmod	85.71
discourse	84.82	discourse	55.56
case	99.42	nsubj	75.00
aux	91.76	case	94.74
cop	97.14	nmod	77.78
vocative	90.38	aux	94.12
iobj	80.43	det	76.92
expl	78.95	vocative	50.00
cc	90.41	cc	100.00
conj	82.35	amod	85.71
obl	75.86	expl	33.33

Table 3:  $IAA$  (F1 calculations used as an inter-annotator agreement metric) for selected high-frequency relations, sorted by frequency separately for CDS and for CS; punctuation excluded.

### 4.2.1 Clitic doubling

Clitic doubling is exhibited by several languages in the Balkan Sprachbund (Tomić, 2011). It is the phenomenon where a clitic pronoun and a full nominal argument co-occur, which makes it a contentious case for UD annotation, as there are two forms with the same function (e.g. *obj*). There are limited examples of clitic doubling in UD-BTB, while the phenomenon is a lot more common in spoken Bulgarian, and therefore in CHILDES-BG. In UD-BTB the annotation of clitic doubling is resolved with the full nominal argument tagged with relation *obj* and the clitic pronominal with relation *expl*, sharing the same head.

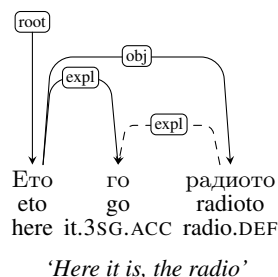


Figure 2: Clitic doubling: the adopted annotation approach. The dashed arc shows an annotation which was decided against.

The guidance all annotators agreed on is in com-

pliance<sup>5</sup> with UD-BTB, and is as follows: the clitic should be attached to the root of the verb (or clause), i.e. to the same head as the nominal argument and be tagged with relation **expl**, see Figure 2. An alternative analysis was proposed (dashed line in Figure 2) where the clitic would have as head the nominal argument itself, because the clitic refers to the nominal argument. This suggestion was ultimately rejected as the nominal argument and the clitic do not always form a constituent.

#### 4.2.2 Vocative case

The case system in Bulgarian has been replaced mostly with prepositions, however the vocative case is still active. The vocative case is used to address someone directly, which makes it especially common in spoken language, and less so in written form, as it requires direct speech between interlocutors. UD has a relation **vocative** especially for this case and this is what we use for annotation (see Figure 3). Sometimes the vocative noun also has the semantic role of a subject. Following a discussion of whether the vocative noun should be introduced via **nsubj** or **vocative** in such cases, it was agreed upon that **vocative** has precedence. This is justified, as Bulgarian is *pro*-drop, so the more likely role of the vocative noun is the vocative role.

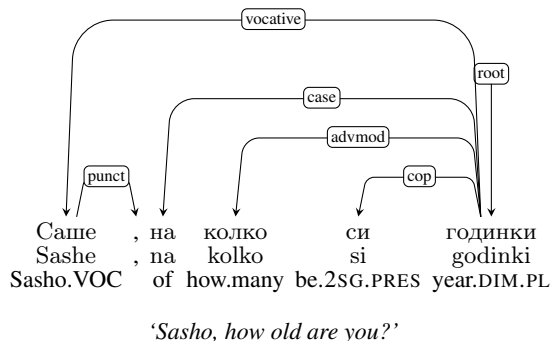


Figure 3: Vocative case.

**Special case: address inversion** is a case of nominal address, where the speaker (addresser) uses their role in the dyad to address the addressee (Beyrer and Kostov, 1978; Braun, 1988; Pavlova, 2015). Most commonly in CHILDES-BG this is exhibited with the vocative of mum: МАМО (*mamo* mum.VOC), however it is also possible with other kinship terms and names. In such cases,

<sup>5</sup>One reviewer suggested that the dislocated relation might be appropriate to use here, however, it is not used at all for Bulgarian. UD-BTB has existing albeit few examples of clitic doubling, hence we stick with the convention.

the consensus among the annotators is to use the vocative relation (see Figure 4).

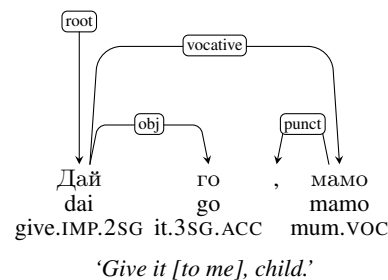


Figure 4: Address inversion.

#### 4.2.3 Reflexive pronoun се/си (se/si refl)

In Bulgarian the form се/си (*se/si* refl) can manifest as three types of reflexive pronoun:<sup>6</sup> dative reflexive, accusative reflexive, and possessive reflexive. Below are examples of these three cases and the appropriate use of the UD relations to annotate each case. Refer to Penčev (1996) for a detailed grammatical account of these forms and to Slavcheva (2003) for the implications of these forms on morphosyntactic annotation.

**Dative reflexive clitic** When си (*si* refl.DAT) is a reflexive dative clitic (*dativus ethicus*), it is introduced with **expl** relation and headed by the verb it relates to (see Figure 5). A *dativus ethicus* is an optional use of the clitic which contributes emotional context (in this context, comfort or indulgence).

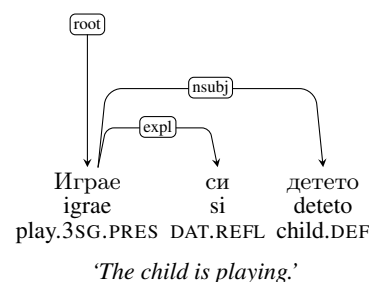


Figure 5: Dative reflexive clitic.

**Accusative reflexive clitic** When си (*si* refl) is an accusative reflexive clitic (*true reflexive*), it is introduced with **expl** relation and headed by the verb it is part of (see Figure 6). The role of the accusative reflexive clitic is to refer to the object, when it is identical to the subject. Verbs which permit or require the use of a true reflexive have that information encoded in their lemmas (not pictured). While the accusative reflexive refers to the semantic

<sup>6</sup>Additionally, it is also the 2SG.PRES form of the verb 'be' as in Figure 3.

object of the sentence predicate, the convention of UD-BTB is to use the `expl` relation rather than `obj`. Using `expl` for reflexive pronouns is an established UD practice for Slavic languages.<sup>7</sup>

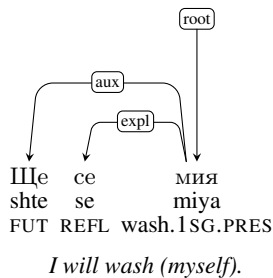


Figure 6: Accusative reflexive clitic.

**Short form of possessive pronoun** If `си` (*si*) can be replaced with a full form possessive pronoun, its relation is `det` with head the nominal which it specifies (see Figure 7), following the case of non-contracted possessive forms.<sup>8</sup> These cases can be ambiguous as the use of `си` (*si*) can resemble a dative ethicus. The annotator agreement is to deterministically apply the rule above when annotating to ensure consistency. Note that this may result in a non-projective parse.

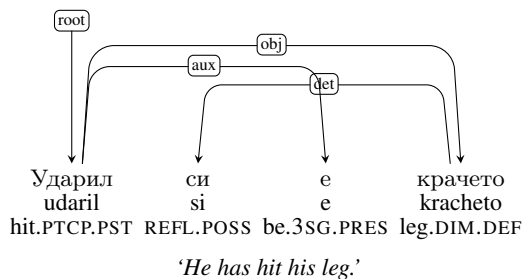


Figure 7: Possessive reflexive clitic.

## 5 Evaluation of the automatic UD parses of CS and CDS

We evaluate the performance of the Stanza parser (Qi et al., 2020) on the manually corrected portion of UD-CHILDES-BG and present key metrics in Table 4. Additionally, we provide existing benchmarks for Bulgarian UD parsing, all based on the test set of UD-BTB: the state-of-the-art Bulgarian UD parser results are from Hromei et al. (2024) and the Stanza results are from the

<sup>7</sup><https://universaldependencies.org/u/dep/expl.html>

<sup>8</sup><https://universaldependencies.org/bg/dep/det.html>

StanfordNLP website.<sup>9</sup> For further review of UD parsers for Bulgarian see Atanasov (2024). As the existing Stanza parser is trained on adult written Bulgarian, UD-BTB (Osenova and Simov, 2017), we expect that it will underperform on relations more prominent in developmental Bulgarian.

System	UPOS	Lemma	UAS	LAS
<i>Benchmark results on UD-BTB from the literature</i>				
U-DepPLLaMA	–	–	96.37	93.77
UDPipe 2.0++	–	–	95.34	92.62
Stanza	98.68	97.29	93.37	90.21
<i>Evaluation of Stanza on Bulgarian CS and CDS</i>				
CDS	92.26	89.52	92.70	85.54
CS	85.60	83.54	90.97	81.52
Overall	89.93	87.43	92.10	84.14

Table 4: Dependency parsing performance for Bulgarian. Benchmark scores are based on training on the UD-BTB: UDPipe 2.0++ and U-DepPLLaMA are from Hromei et al. (2024) and the Stanza scores are from the StanfordNLP website. In this work we evaluate the Stanza parser against manually annotated Bulgarian CDS and CS.

A per-relation breakdown of the Stanza parser is reported in Table 5. As expected, the parser performs worse on child speech for the majority of the relations. Core syntactic relations such as case, aux, and cop achieve very reliable F1. Relations more common in spoken language such as discourse and vocative are relatively frequent in the corpus (5.9% and 3.3% of all dependencies respectively), yet they exhibit the lowest parsing accuracy, which highlights that CS and CDS exhibit phenomena less prominent in written adult language, of which UD-BTB is comprised.

### 5.1 Inability to handle vocative case

Although the use of the vocative case is a relatively frequent phenomenon across both CS and CDS (3.3% of all dependencies), the Stanza parser has lowest UAS and LAS on the vocative relation. A case study based on the most common vocative in CHILDES-BG, `МАМО` (*mamo* mum.VOC), shows that only 7/1,848 of its occurrences in the CDS part of CHILDES-BG are correctly labelled with relation vocative. The rest of the time `МАМО` (*mamo* mum.VOC) is tagged as an interjection (INTJ) and introduced with relation discourse instead. Most surprisingly, 598 times `МАМО` (*mamo*

<sup>9</sup><https://stanfordnlp.github.io/stanza/performance.html>

Relation	% of total	CDS F1	CS F1	Overall F1
case	4.03%	99.50	94.01	97.89
aux	5.16%	97.20	92.15	96.19
root	30.22%	94.07	94.33	94.19
cop	3.32%	95.24	94.79	95.14
cc	2.18%	94.62	85.29	91.78
advmod	7.34%	95.03	87.85	93.10
<i>Stanza UD-BTB mean LAS=90.21</i>				
obj	9.68%	89.07	80.04	86.22
amod	1.18%	89.12	86.79	88.07
ccomp	1.98%	81.48	73.76	79.67
nsubj	7.20%	85.83	78.25	83.47
det	1.91%	85.71	83.24	84.87
conj	5.45%	76.90	83.52	81.19
expl	2.16%	78.73	76.47	78.10
iobj	2.38%	75.89	73.08	75.23
discourse	6.08%	78.80	54.83	73.02
xcomp	0.71%	71.95	39.02	65.37
obl	1.86%	65.02	55.17	61.68
advcl	0.67%	69.57	44.83	63.01
nmod	1.15%	64.42	49.16	57.36
vocative	3.16%	28.09	24.88	26.93

Table 5: Dependency relation performance (F1) for Bulgarian child-directed speech (CDS), child speech (CS), and the combined dataset. Relations are grouped into high- and low-performing categories based on overall F1. Percentages indicate the proportion of all gold dependency relations across the CS and CDS annotated dataset.

mum.VOC), which is not a homonymic form, is mistakenly labelled as a verb with lemma МАМ-(се) (*mam-(se)*), which cannot be found in dictionaries for Bulgarian or in the UD-BTB corpus. Interpreting МАМО (*mamo* mum.VOC) as a verb is ungrounded linguistically.

## 5.2 Reflexive pronoun СИ (si refl)

The clitic СИ (*si* refl) poses a challenge to the parser as it can manifest as various types of reflexive pronouns. The relevant dependency relations to СИ (*si* refl) are *expl* used for the accusative and dative reflexive pronoun, and *det* used for the short form of the possessive reflexive (see subsection 4.2.3).

In order to correctly establish the dependency relation, the parser first needs to predict the morphological information of СИ (*si*). Although the parser often successfully distinguishes between the accusative/dative and the possessive (which can be discerned by the correct assignment of the lemma as the long form possessive pronoun), in the possessive case it does not correctly attach it to the object with the *det* relation, despite having the capability to create non-projective parses (see Figure 7).

## 5.3 Child speech

As expected, the parser performs worse on CS than on CDS (see Table 4 and Table 5). CS may contain utterances which are ungrammatical to adult speakers, due to atypical word order and/or omitted function words such as case markers, which the parser has no way of recovering. The below CS sentence, Figure 8, exhibits a non-standard word order as well as an omitted function word. Although the typical order in Bulgarian is SVO, relatively free word order is allowed and subject drop is allowed. The order in Figure 8 is IO (indirect object) - V (verb with implied subject) - O (object), which is not grammatical for adults, but it is still recoverable. The parser wrongly assigns the verb to be in third person, and due to the positioning of the indirect object at what would typically be the overt subject position, wrongly assigns the indirect object as a subject. The example is made more challenging as there is an omitted case marker на (*na* ‘on’), which the annotators can recover, but the parser, which relies purely on the surface form, cannot.

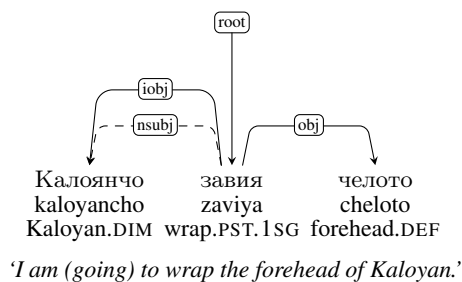


Figure 8: Child speech with non-standard word order. Incorrect relation assignment by the parser on child speech, *nsubj*, is illustrated with a dashed arc. The correct relation is *iobj*, as assigned by an annotator.

## 6 Conversion of dependency to constituency parses

Dependency parses can be converted to constituency parses following the Collins method (Collins et al., 1999), if constituency parses are required or preferred. The dependency to constituency conversion based on the Collins method does not provide support for non-projective dependency parses. Non-projective parses comprise 1.45% of the manually corrected sentences. We tested the conversion algorithm from Kando et al. (2022), which is available on GitHub,<sup>10</sup> and augmented it to

<sup>10</sup><https://github.com/gifdog97/dep-to-const>

ensure the constituency grammar was more linguistically meaningful. Specifically, we added a distinction between pre-terminals and non-terminals and added a start symbol ROOT for all utterances. This facilitates the use of the resulting constituency parses with common NLP libraries such as NLTK. The automatically converted constituency parses will be made available alongside the publication of the rest of the treebank and the augmentation of the conversion algorithm will be made available on GitHub upon publication.

## 7 Conclusion

This paper presents (i) UD-CHILDES-BG, a manually corrected Universal Dependencies (UD) treebank of Bulgarian child and child-directed speech, (ii) a quantitative evaluation of inter-annotator agreement on developmental data, and (iii) a systematic analysis of parser errors on developmental Bulgarian. UD-CHILDES-BG is comprised of 10% manually corrected parses, of which 14% are independently double-annotated. We find that parsing errors arise mainly from properties of Bulgarian discourse and child-specific constructions. The resulting resource constitutes the first manually corrected CS and CDS treebank for a Balkan or Slavic language and offers a methodological account of annotation correction and error analysis in linguistically challenging data.

## Limitations

The data from CHILDES-BG contains some inconsistencies: words are not always canonically transcribed and the CHILDES CHAT annotation conventions are not always followed consistently. For the manually corrected part of UD-CHILDES-BG, after automatic preprocessing of the CHAT annotations, the utterances were further manually normalised. INCEpTION is a suitable tool for this project as it requires no installation and can be accessed online by all annotators. However, it has limited support for the morphological features layer, which could have negatively affected the morphological feature annotation.

An additional limitation of this project and a direction for future work is the retraining of Bulgarian UD parsers using the manually corrected portion of UD-CHILDES-BG as supplementary training data. Given the performance decrease observed on CS and CDS relative to adult written Bulgarian, fine-tuning on developmental and spoken language may

improve parsing accuracy for discourse-related and child-specific constructions. Extending the manually corrected portion of the corpus beyond the current 10% sample would also support further investigation of parser adaptation at larger scale.

## Acknowledgments

We would like to thank Professor Svetla Koeva for connecting the annotators with each other.

## References

- Elena Andonova. 2015. [Parental report evidence for toddlers’ grammar and vocabulary in bulgarian](#). *First Language*, 35(2):126–136.
- Atanas Atanasov. 2024. [Dependency parser for Bulgarian](#). In *Proceedings of the Sixth International Conference on Computational Linguistics in Bulgaria (CLIB 2024)*, pages 98–105, Sofia, Bulgaria. Department of Computational Linguistics, Institute for Bulgarian Language, Bulgarian Academy of Sciences.
- A. Beyrer and K. Kostov. 1978. “Umgekehrte Anrede” im Bulgarischen und Rumänischen? (“Address inversion” in Bulgarian and Rumanian?). *Balkansko Ezikoznanie*, 21(4):41–53.
- Núria Bosch. 2025. [Categorial granularity in syntactic acquisition: A multilingual corpus study on the left periphery](#). *Glossa*, 10.
- Friederike Braun. 1988. *Terms of Address: Problems of Patterns and Usage in Various Languages and Cultures*, volume 50 of *Contributions to the Sociology of Language*. Mouton de Gruyter, Berlin and New York and Amsterdam.
- Alastair Butler, Susanne Miyata, and Yumiko Kinjo. 2022. The Soyogo Treebank – a parsed corpus of child Japanese. <https://soyogo.github.io>.
- Michael Collins, Jan Hajic, Lance Ramshaw, and Christoph Tillmann. 1999. [A statistical parser for Czech](#). In *Proceedings of the 37th Annual Meeting of the Association for Computational Linguistics*, pages 505–512, College Park, Maryland, USA. Association for Computational Linguistics.
- Marie-Catherine de Marneffe, Christopher D. Manning, Joakim Nivre, and Daniel Zeman. 2021. [Universal dependencies](#). *Computational Linguistics*, page 1–54.
- Shai Gretz, Alon Itai, Brian MacWhinney, Bracha Nir, and Shuly Wintner. 2013. [Parsing hebrew childes transcripts](#). *Language Resources and Evaluation*, 49(1):107–145.
- Arzu Burcu Güven, Anna Rogers, and Rob Van Der Goot. 2025. [Do syntactic categories help in developmentally motivated curriculum learning for language](#)

- models? In *Proceedings of the First BabyLM Workshop*, pages 288–300, Suzhou, China. Association for Computational Linguistics.
- Claudiu Daniel Hromei, Danilo Croce, and Roberto Basili. 2024. **U-DepPLLaMA: Universal dependency parsing via auto-regressive large language models**. *Italian Journal of Computational Linguistics*, 10(1).
- Shunsuke Kando, Hiroshi Noji, and Yusuke Miyao. 2022. **Multilingual syntax-aware language modeling through dependency tree conversion**. In *Proceedings of the Sixth Workshop on Structured Prediction for NLP*, pages 1–10, Dublin, Ireland. Association for Computational Linguistics.
- Vera Kempe and Patricia J. Brooks. 2001. **The role of diminutives in the acquisition of russian gender: Can elements of child-directed speech aid in learning morphology?** *Language Learning*, 51(2):221–256.
- Jan-Christoph Klie, Michael Bugert, Beto Boullosa, Richard Eckart de Castilho, and Iryna Gurevych. 2018. **The INCEpTION platform: Machine-assisted and knowledge-oriented interactive annotation**. In *Proceedings of the 27th International Conference on Computational Linguistics: System Demonstrations*, pages 5–9, Santa Fe, New Mexico. Association for Computational Linguistics.
- Jan-Christoph Klie, Bonnie Webber, and Iryna Gurevych. 2023. **Annotation error detection: Analyzing the past and present for a more coherent future**. *Computational Linguistics*, 49(1):157–198.
- Houjun Liu. 2024. **Morphosyntactic analysis for CHILDES**. *Language Development Research: An Open-Science Journal*, 4.
- Zoey Liu and Emily Prud’hommeaux. 2023. **Data-driven parsing evaluation for child-parent interactions**. *Transactions of the Association for Computational Linguistics*, 11:1734–1753.
- Brian MacWhinney. 1992. **The CHILDES project: tools for analyzing talk**. *Child Language Teaching and Therapy*, 8(2):217–218.
- Brian MacWhinney. 2012. **Morphosyntactic analysis of the CHILDES and TalkBank corpora**. In *Proceedings of the Eighth International Conference on Language Resources and Evaluation (LREC’12)*, pages 2375–2380, Istanbul, Turkey. European Language Resources Association (ELRA).
- Jan Odijk, Alexis Dimitriadis, Martijn van der Klis, Marjo van Koppen, Meie Otten, and Remco van der Veen. 2018. **The AnnCor CHILDES treebank**. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan. European Language Resources Association (ELRA).
- Petya Osenova and Kiril Simov. 2004. **BTB-TR05: BulTreeBank Stylebook**. Technical report, Bulgarian Academy of Sciences.
- Petya Osenova and Kiril Simov. 2017. **Recent developments within BulTreeBank**. In *Proceedings of the 16th International Workshop on Treebanks and Linguistic Theories*, pages 129–137, Prague, Czech Republic.
- Neda Pavlova. 2015. **“Reverse addresses” in Bulgarian speech – between kinship appellatives and discursive markers**. *Balkanistic Forum*, 21(1).
- Lisa S. Pearl and Jon Sprouse. 2013. **Syntactic islands and learning biases: Combining experimental syntax and computational modeling to investigate the language acquisition problem**. *Language Acquisition*.
- Jordan Penčev. 1996. **Functions of the formant se/si in bulgarian**. *Revue des études slaves*, 68(4):497–515.
- Slav Petrov, Leon Barrett, Romain Thibaux, and Dan Klein. 2006. **Learning accurate, compact, and interpretable tree annotation**. In *Proceedings of the 21st International Conference on Computational Linguistics and 44th Annual Meeting of the Association for Computational Linguistics*, pages 433–440, Sydney, Australia. Association for Computational Linguistics.
- Carl Pollard and Ivan A. Sag. 1994. *Head-Driven Phrase Structure Grammar*. The University of Chicago Press, Chicago.
- Velka Popova. 2023. **The Emergence of Word Classes in Early Bulgarian Language Ontogenesis. A Pilot Corpus Study**. *Journal of Bulgarian Language*, 70(PRIL):290–308.
- Velka Popova and Vladimir Filipov. 2022. **Acquisition of Modal Verbs in Bulgarian: Analysis of Longitudinal Data from CHILDES**. *Ezikov Svyat (Orbis Linguarum)*, (ezs.swu.v20i3):332–346.
- Velka Popova and Dimitar Popov. 2020. **CHILDES Bulgarian LabLing Corpus**.
- Peng Qi, Yuhao Zhang, Yuhui Zhang, Jason Bolton, and Christopher D. Manning. 2020. **Stanza: A python natural language processing toolkit for many human languages**. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 101–108, Online. Association for Computational Linguistics.
- Kenji Sagae, Eric Davis, Alon Lavie, Brian MacWhinney, and Shuly Wintner. 2010. **Morphosyntactic annotation of CHILDES transcripts**. *Journal of Child Language*, 37(3):705–729.
- Kenji Sagae, Alon Lavie, and Brian MacWhinney. 2001. **Parsing the CHILDES database: Methodology and lessons learned**. In *Proceedings of the Seventh International Workshop on Parsing Technologies*, pages 166–176, Beijing, China.
- Kiril Simov and Petya Osenova. 2003. **Practical Annotation Scheme for an HPSG Treebank of Bulgarian**. In *Proceedings of the 4th International Workshop on Linguistically Interpreted Corpora (LINC-2003)*, pages 17–24.

- Kiril Simov, Petya Osenova, Alexander Simov, and Milen Kouylekov. 2005. Design and Implementation of the Bulgarian HPSG-based Treebank. *Journal of Research on Language and Computation. Special Issue*, pages 495–522.
- Kiril Simov, Petya Osenova, and Milena Slavcheva. 2004. [BTB-TR03: BulTreeBank Morphosyntactic Tagset](#). Technical report.
- Kiril Simov, Petya Osenova, Milena Slavcheva, Sia Kolkovska, Elisaveta Balabanova, Dimitar Doikoff, Krassimira Ivanova, Alexander Simov, and Milen Kouylekov. 2002a. [Building a linguistically interpreted corpus of Bulgarian: the BulTreeBank](#). In *Proceedings of the Third International Conference on Language Resources and Evaluation (LREC'02)*, Las Palmas, Canary Islands - Spain. European Language Resources Association (ELRA).
- Kiril Simov, Gergana Popova, and Petya Osenova. 2002b. HPSG-based syntactic treebank of Bulgarian (BulTreeBank). In Paul Rayson Andrew Wilson and Tony McEnery, editors, *A Rainbow of Corpora: Corpus Linguistics and the Languages of the World*, pages 135–142. Lincom-Europa.
- Milena Slavcheva. 2003. [Some aspects of the morphological processing of Bulgarian](#). In *Proceedings of the 2003 EACL Workshop on Morphological Processing of Slavic Languages*, pages 71–77, Budapest, Hungary. Association for Computational Linguistics.
- Ida Szubert, Omri Abend, Nathan Schneider, Samuel Gibbon, Louis Mahon, Sharon Goldwater, and Mark Steedman. 2024. [Cross-linguistically consistent semantic and syntactic annotation of child-directed speech](#). *Language Resources and Evaluation*, 59(2):727–776.
- Luka Terčon and Nikola Ljubešić. 2023. [CLASSLA-Stanza: The Next Step for Linguistic Processing of South Slavic Languages](#). *arXiv preprint*.
- Olga Mišeska Tomić. 2011. *16 Balkan Sprachbund features*, page 307–324. DE GRUYTER MOUTON.
- Leon Weber-Genzel, Siyao Peng, Marie-Catherine De Marneffe, and Barbara Plank. 2024. [VariErr NLI: Separating annotation error from human label variation](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2256–2269, Bangkok, Thailand. Association for Computational Linguistics.
- Xiulin Yang, Zhuoxuan Ju, Lanni Bu, Zoey Liu, and Nathan Schneider. 2025. [UD-English-CHILDES: A collected resource of gold and silver Universal Dependencies trees for child language interactions](#). In *Proceedings of the Eighth Workshop on Universal Dependencies (UDW, SyntaxFest 2025)*, pages 52–58, Ljubljana, Slovenia. Association for Computational Linguistics.

## A Sentence assignment to annotators

### A1 (MM)

- CS: ALE\_cs.conllu (sentences 66-245; 321-379)
- CDS: ALE\_cds.conllu (sentences 1-215; 321-430)
- CS: BOG\_cs.conllu (sentences 1-97)
- CDS: BOG\_cds.conllu (sentences 1-148)
- CDS: TEF\_cds.conllu (sentences 1-10)
- CS: SIM\_cs.conllu (sentences 101-400; 505-924)
- CDS: SIM\_cds.conllu (sentences 161-805)
- CS: ELI\_cs.conllu (sentences 1-88)
- CDS: ELI\_cds.conllu (sentences 101-400; 505-924)

### A2 (YV)

- CS: BOG\_cs.conllu (sentences 1-92)
- CDS: BOG\_cds.conllu (sentences 1-142)
- CS: TEF\_cs.conllu (sentences 276-336)
- CDS: TEF\_cds.conllu (sentences 266-392)

### A3 (YC)

- CS: ALE\_cs.conllu (sentences 1-65)
- CDS: ALE\_cds.conllu (sentences 1-90)
- CS: ALE\_cs.conllu (sentences 246-320)
- CDS: ALE\_cds.conllu (sentences 216-320)

### A4 (TK)

- CS: SIM\_cs.conllu (sentences 1-100)
- CDS: SIM\_cds.conllu (sentences 1-160)
- CS: SIM\_cs.conllu (sentences 401-515)
- CDS: SIM\_cds.conllu (sentences 161-335)

### A5 (IP)

- CS: TEF\_cs.conllu (sentences 1-65)
- CDS: TEF\_cds.conllu (sentences 1-90)

- CS: TEF\_cs.conllu (sentences 65-180)

- CDS: TEF\_cds.conllu (sentences 90-265)

### A6 (TS)

- CS: ELI\_cs.conllu (sentences 1-100)
- CDS: ELI\_cds.conllu (sentences 1-160)
- CS: ELI\_cs.conllu (sentences 101-121)
- CDS: ELI\_cds.conllu (sentences 386-560)
- CS: TEF\_cs.conllu (sentences 181-275)

## B Annotation Scheme

### B.1 Scope

The starting point is automatically generated Universal Dependencies (UD) parses produced using Stanza. For every sentence, annotators perform the following checks (and corrections if necessary), in order:

- Lemma
- UPOS
- Dependency head
- Dependency relation
- Morphological features (when clearly recoverable)
- XPOS tags are **not** corrected

The lemmas, UPOS tags, dependency edges, dependency labels, and morphological features should follow the guidelines of the UD Bulgarian dataset. An exhaustive list of the UPOS tags and dependency tags are explained with examples here: [https://universaldependencies.org/treebanks/bg\\_btb/index.html](https://universaldependencies.org/treebanks/bg_btb/index.html)

Gold labelled UD treebank for Bulgarian based on BulTreeBank: [https://github.com/UniversalDependencies/UD\\_Bulgarian-BTB](https://github.com/UniversalDependencies/UD_Bulgarian-BTB) (files: `bg_btb-ud-dev.conllu`, `bg_btb-ud-test.conllu`, `bg_btb-ud-train.conllu`)

## B.2 Step-by-Step Annotation Procedure

For each sentence:

1. Read the entire utterance.
2. Correct token-level information (lemma and UPOS).
3. Correct dependency heads.
4. Correct dependency relations.
5. Review morphological features when present.

## B.3 Lemma Correction

- Lemmas must use standard Bulgarian orthography (dictionary entries) and be lowercased.
- The lemma of a proper name should also be lowercased, e.g. the lemma of Ани is ани).
- Normalise obvious child and regional variation (e.g. the lemma of икам is искам; the lemma of ши is ще; Кутийката-а-а-а – кутийка; панна [падна] – падам).
- Use the infinitive/base form for verbs (1st person singular, present).
- For verbs with reflexive clitic the lemma is скрия-(се), don't remove the dash, it is convention.
- If the surface form is diminutive, keep the lemma diminutive (e.g. книжката – книжка).
- Use nominative for the lemma of vocative nouns (e.g. the lemma for мама is мама)
- The lemma of тати is unclear whether it should be тати or тате so for consistency we keep тати.
- For pronouns and function words, use standard dictionary lemmas
- For pronouns use 1st person singular, following the gold standard, e.g. ѝ – аз; го – аз, си – се or си – свой.
- xxx is only used in CHILDES when the surface form is not recoverable. If the surface form is xxx, the lemma is also xxx.
- If the surface form is an unrecognisable word, its lemma is the lowercased version of it, e.g. Тпру → тпру.

- Repetitions: single lemma: Дай-дай-дай-дай → дай
- Example of frequent child words and their adult pairings, which should be used as their lemmas if recoverable: йейя – леля, тетито/тефи/тефани/тефче – стефи/стефче/стефани, икам – искам, а – на,

## B.4 UPOS Correction

Ensure UPOS matches syntactic behaviour. Use UPOS=X only when no category is recoverable.

- бе, хайде, айде, де all have UPOS tag interjection INTJ
- ще is AUX
- какво is tagged as DET in the gold
- да is AUX when used in verbal constructions, and PART when used in discourse
- Tokens marked xxx: recover UPOS if possible; otherwise use UPOS=X.
- Proper nouns: may not always be capitalised as expected, but treat them as PROP (e.g. калоянчо, мончи)

## B.5 Morphological Features

Do not guess features if the form is unclear. Leave empty.

### Verbal

Aspect=Imp|Mood=Ind|Number=Sing|Person=3|Tense=Pres|VerbForm=Fin|Voice=Act

- Mood: Cnd, Imp, Ind
- Number: Plur, Sing
- Tense: Imp, Past, Pres
- Person: 1, 2, 3
- Aspect: Imp, Perf
- Voice: Act, Pass
- VerbForm: Fin, Part

### Nominal

Definite=Def|Gender=Fem|Number=Sing

- Gender: Fem, Masc, Neut
- Case: Voc

- Definite: Def, Ind
- Number: Plur, Sing

Notes on the features:

- We only use Case=Voc for nouns but not for verbs.
- Mark vocative case even when the vocative form is the same as the nominal form (e.g. вуйчо, дай ми го).
- Definite: even if not with **ЪТ** correctly when it is in a role of a subject, e.g. **КЪЛВАЧА**, the morphological feature should still include Definite=Def.

## B.6 Punctuation

- Attach punctuation using **punct**.
- Attach to the syntactic head of the phrase that necessitates the use of the punctuation (not always the sentence root).

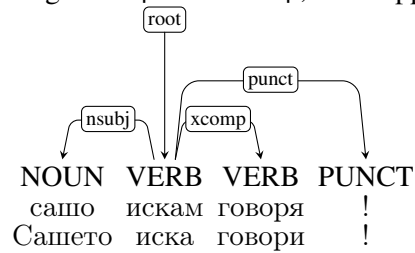
## B.7 Dependency Head Correction

**IMPORTANT:** In INCEpTION dependency arcs are drawn from the UPOS box (blue) of the head to the UPOS box of the child. Hold Shift, select the UPOS of the head (From), and then select the POS of the dependent (To), let go of Shift. Edit the dependence relation.

Checklist:

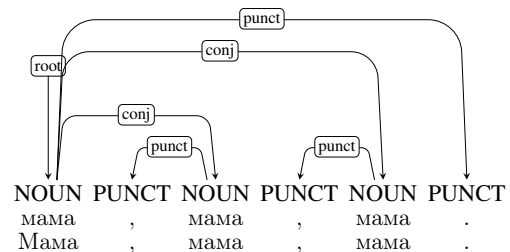
- Exactly one root per sentence.
- Verify attachment directionality: the arrow should be pointing from the head to the child.
- Finite verb is typically the root.
- In copular clauses, the predicate nominal is the root (**not** the subject).
- In copula questions with k-word, the k-word is usually the root.
- In fragments, the semantic head is the root.
- Single-word utterances receive root relation.
- In cases of repetition, first token is root; others attach via **conj**.
- When no element is subordinated (e.g. Христo Смирненски the leftmost element is the root).

- In child utterances where **да** is omitted, attach infinitival verb as in adult speech, using **xcomp** or **ccomp**, as appropriate.



- Coordination:

- First conjunct attaches to the head.
- Subsequent conjuncts attach via **conj**.
- Coordinating conjunction, if present, attaches via **cc**.



- Punctuation attach to the particle/clause that necessitates it – not always the root (see above).
- Adposition attachment to the head of the phrase rather than to interrogative modifier.

## B.8 Dependency Relation Correction

Common distinctions (review the examples file UD\_relations\_example\_for\_each\_relation on INCEpTION). Below is a list of all relations and how to distinguish between them:

- **xcomp vs ccomp:**

- **xcomp** = embedded clause without its own subject (the subject of the xcomp is inherited from the higher clause);
- **ccomp** = embedded clause with its own subject.

- **parataxis vs conj** In UD-BulTreeBank, parataxis is used exclusively for direct speech in a sentence. Following the gold standards, when there are clauses which are equal (neither is subordinated) we use **conj** even if a **cc** is missing.

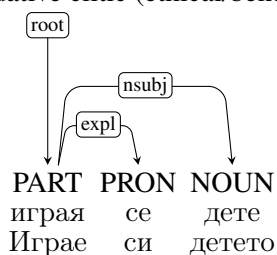
- **nmod vs obl:**
  - nmod nominal modifier modifies a noun or nominal phrase.
  - obl oblique modifies a verb/verbal phrase; it introduces **optional** circumstantial information.
- **iobj vs obl:** not everything with an adposition is indirect object; consider the verb valency and whether the sentence would be complete without this phrase – if the phrase is optional, it is obl.
- **fixed vs flat:**
  - fixed = rigid multiword expressions with no internal syntax (e.g. multi-word coordinating conjunction, such as само че);
  - flat = names or sequences with internal structure suppressed (e.g. proper names, such as Кума Лиса).
- **nmod vs flat:** In noun+name expressions (e.g. баба Лили лили has head баба with relation nmod (as opposed to flat; flat is for e.g. name + surname)
- **expl (expletive):** Reflexive pronoun e.g. се, си; as well as the clitic in clitic doubling.
- **vocative:** Direct address to interlocutor (e.g. мамо).
- **discourse:** Discourse particles or interjections (e.g. бе, хайде, де, ами).
- **root:** Main predicate of the sentence.
- **amod:** Adjectival modifier of a noun.
- **advmod:** Adverbial modifier of verb/adjective.
- **det:** Determiner modifying a noun; can also be used for possessive clitics (e.g. крачето си, неговия учител, as well as for numerical determiner един момък дойде; една жена ми каза
- **case:** Adposition introducing nominal dependents.
- **cc:** Coordinating conjunction.
- **conj:** Conjunct in coordination.

- **cop:** Copula (usually forms of съм).
- **aux / aux:pass:** Auxiliary verbs; passive auxiliaries use aux:pass.
- **nsubj / nsubj:pass:** Nominal subject (active / passive).
- **obj:** Direct object.
- **iobj:** Indirect object.
- **mark:** Subordinating marker (e.g. че)
- **acl / acl:relcl:** Clausal modifier of noun; for relative clauses introduced by a k-word (който/която/които), use acl:relcl, otherwise use acl
- **advcl:** Adverbial clause modifier.
- **advcl vs obl:** advcl is for clauses (with verb), whereas obl is for noun phrases (for преди да отиде use advcl, for преди обед use obl).
- **appos:** Appositional noun phrases.
- **nummod:** Numerical modifier.
- **punct:** Punctuation.

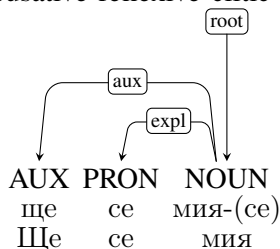
## B.9 Common difficult cases with examples

- Expletive си vs determiner си

- **expl** if it is part of the verb, as reflexive dative clitic (ethical/benefactive dative)

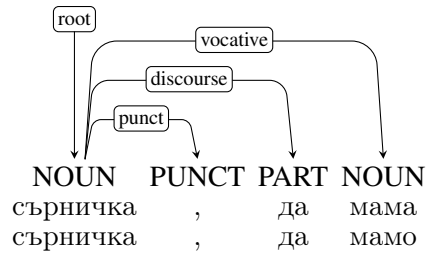
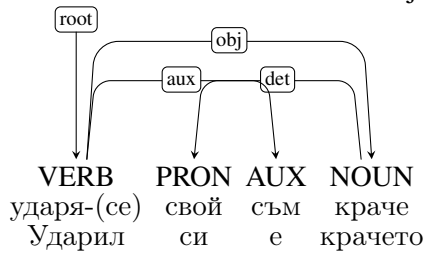


- **expl** if it is part of the verb, as accusative reflexive clitic (true reflexive)

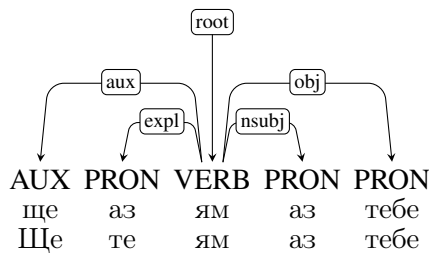
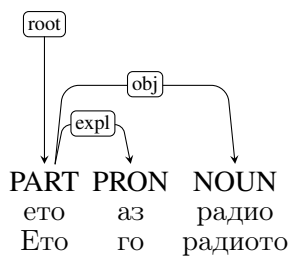


- **det** if it can be replaced with свой (if it can be replaced with свой/своя/свои,

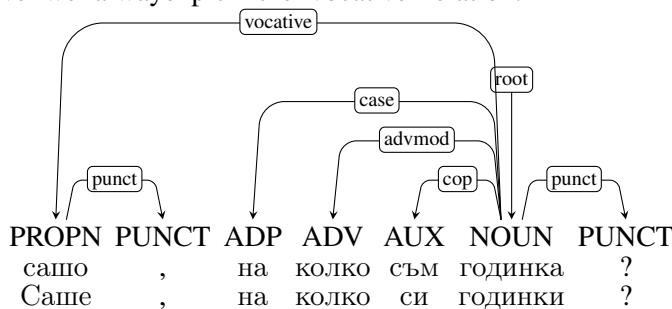
then its lemma is **свой** and the relation is **det** with head the object)



- Clitic doubling: the clitic should be attached to the root of the verb (or clause), i.e. to the same head as the object and be tagged with relation **expl**



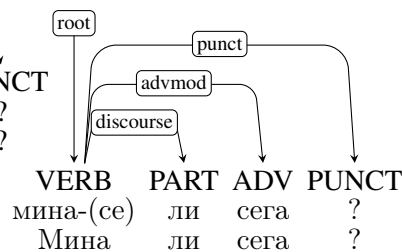
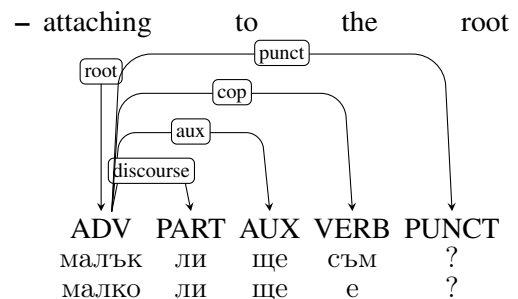
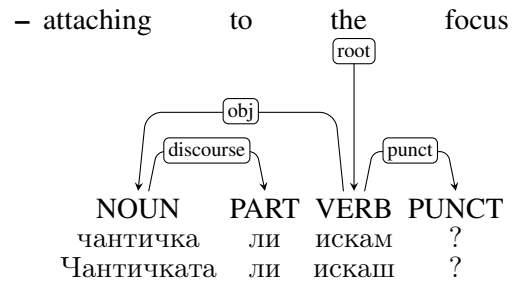
- nsubj or vocative: When a form could be analysed as both nsubj and vocative we always pick the vocative relation.



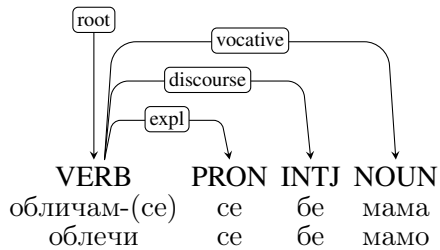
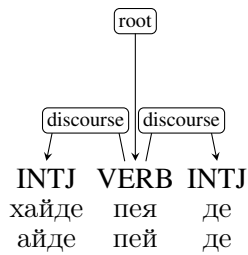
- Attach punctuation to the head that necessitates it – e.g. in the above example, Саше, на колко си годинки? the vocative noun necessitates the comma, so it is attached to it.
- Attach adposition to the same head as the head of the k-word.

- **мамо** referring to the child: treat it as a standard vocative noun (not as adjective **мамин**)

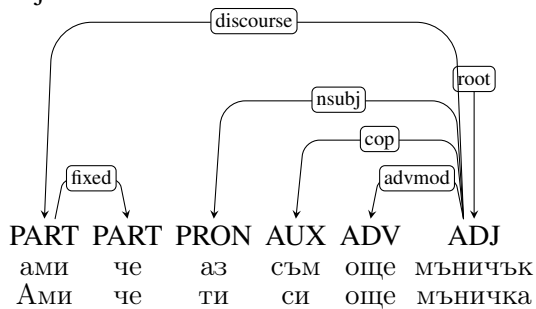
- attachment of **ли**: In the gold standard it is attached to the nearest item. **ли** is a clitic which changes the focus (topic) so attach it to the focus of the sentence, which might be the root. If it is ambiguous, then attach it to the root.



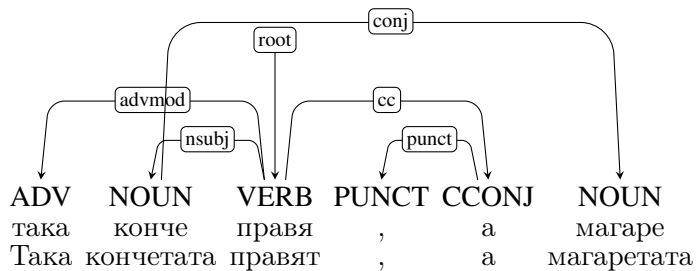
- Multiple interjections (discourse items): test by removing each interjection, and if the sentences with one interjection is still grammatical, then attach the interjections independently to the root using relation **discourse**



- multi-word coordinating conjunctions: use **fixed**



- ellipsis – difficult examples, excluded from the UD BulTreeBank. Attach words to the corresponding word with the same part of speech.



## B.10 Final checks

- Exactly one root
- No orphan nodes

If uncertain:

- Follow the existing gold standard: [https://github.com/UniversalDependencies/UD\\_Bulgarian-BTB](https://github.com/UniversalDependencies/UD_Bulgarian-BTB)
- Leave a comment in the annotation tool.

# IndiAnn: A Web-based Annotation Platform for Indic Languages

Bandaru Lavadeep, Ritwik Raghav, Abhik Jana

IIT Bhubaneswar, India

{22cs01049,a23cs09001,abhikjana}@iitbbs.ac.in

## Abstract

Linguistic annotation tools that work well for non-Indic languages (e.g. English, German, Spanish, etc.) often fail with Indic scripts due to complex Unicode properties, including visual reordering of vowel matras, conjunct characters, and grapheme clusters spanning multiple code points. In this paper, we present a web-based annotation platform *IndiAnn*, designed for low-resource Indic languages, which uses native browser Unicode rendering, offset-based storage that preserves grapheme clusters, and no forced tokenization in the user interface. The tool supports annotation for tasks such as part-of-speech (POS) tagging, named entity recognition (NER), dependency relation annotation, and semantic role labelling (SRL), that maintain correct character boundaries and enable seamless interoperability with standard NLP pipelines and tools. The framework is designed for Indic languages and has been tested on Telugu, Hindi, Tamil, Malayalam, Bengali, Odia, Marathi, and Kannada, with no script breakage during annotation. To the best of our knowledge, this is the first ever attempt at building a unified annotation framework (IndiAnn), which covers annotation for such varieties of key NLP tasks, having provision for eight Indic languages. The code repository is made publicly available<sup>1</sup>.

## 1 Introduction

Linguistic annotation is essential for building high-quality datasets for natural language processing, especially for low-resource languages. Large annotated resources such as Universal Dependencies (Nivre et al., 2016) and shared tasks like CoNLL-2003 (Sang and Meulder, 2003) highlight the importance of structured annotation for downstream NLP tasks. However, many languages, particularly Indic languages, still lack sufficient annotated resources (Joshi et al., 2020; Maji et al., 2025).

Most existing annotation tools are designed primarily for languages such as English, German, Spanish, etc. and assume a simple, linear character model with fixed tokenization. These assumptions do not hold for Indic languages, which have more complex writing systems. Indic scripts present several challenges due to their Unicode properties. In many cases, the storage order of characters differs from their visual display order, particularly because of vowel matras and combining marks (The Unicode Consortium, 2022b; Ansary et al., 2024). Moreover, what appears as a single character to the annotator is often composed of multiple code points forming a grapheme cluster, which is the more appropriate unit for segmentation in these scripts (Ansary et al., 2024). As a result, tools that rely on byte-level or code-point indexing often produce incorrect offsets, leading to incorrect span boundaries, cursor misalignment, and broken annotation highlighting. These issues directly impact annotation quality and make it difficult to build reliable datasets. Widely used annotation tools such as WebAnno (Yimam et al., 2013), INCEPTION (Klie et al., 2018), and Brat (Stenetorp et al., 2012) do not explicitly address these challenges, as they are primarily designed for languages with simpler orthographic structures. In addition, many of these systems rely on fixed tokenization schemes, which can incorrectly split or merge units in Indic languages, especially in the presence of conjunct characters or multi-word expressions.

Existing efforts for Indian languages, such as the Paninian dependency annotation scheme by Begum et al. (2008) and the ILCIANN tool by Kumar et al. (2021), are typically designed for specific tasks and operate over pre-tokenized or structured text. While effective for large-scale annotation and linguistic analysis, they do not support direct span selection over raw text as in tools like WebAnno, and therefore do not explicitly address character-level challenges in Indic scripts, such as grapheme

<sup>1</sup><https://github.com/Lavadeep/INDIANN>

clusters and rendering-order mismatches.

In this paper, we present a web-based annotation platform - **IndiAnn**, designed specifically for Indic languages. The system is built around native Unicode rendering and uses a single canonical text with offset-based annotation that preserves grapheme clusters. Unlike traditional tools, it does not enforce tokenization in the user interface, allowing annotators to select natural linguistic units directly. The proposed system supports multiple annotation layers, including span-based annotations like POS tagging and NER, relations between spans, dependency relation annotation, and semantic role labeling. It accepts input in a variety of formats, including Plain text, Word documents, PDF files, and CoNLL-U, enabling flexible data ingestion. For interoperability, the system provides export functionality in JSON format, as well as CoNLL-U for POS and dependency annotations. The platform is designed for Indic languages and has been tested on Telugu, Hindi, Tamil, Malayalam, Bengali, Odia, Marathi, and Kannada, ensuring compatibility with diverse scripts and linguistic structures.

## 2 Related Work

Several annotation platforms have been developed for linguistic annotation, including WebAnno (Yimam et al., 2013), INCEPTION (Klie et al., 2018), and Brat (Stenetorp et al., 2012). These systems support multiple annotation layers and collaborative workflows, and have been widely used to annotate documents written in European languages. More recent tools such as doccano (Nakayama et al., 2018), and Label Studio (Tkachenko et al., 2020-2025) provide flexible web-based annotation interfaces.

Despite their capabilities, these tools are largely designed with assumptions that align with European languages and simpler writing systems. Prior work has shown that many NLP tools implicitly rely on whitespace-based tokenization and linear character representations, which do not generalize well to languages with complex scripts (Bird et al., 2009). This limitation becomes more pronounced in multilingual settings, where language-specific properties are often not adequately handled (Bender, 2011).

In the context of Indic languages, several efforts have been made to develop resources and tools tailored to local linguistic phenomena. Be-

gum et al. (2008) propose a dependency annotation scheme based on the Paninian framework, introducing syntactic-semantic relations suitable for morphologically rich, free-word-order languages such as Hindi. Their approach operates over pre-tokenized and chunked text, where annotations are defined between predefined units rather than directly on raw text spans. Similarly, Kumar et al. (2021) presents ILCIANN, a web-based annotation tool designed to create and manage large-scale parallel corpora across multiple Indian languages. The system supports distributed annotation and project management, but restricts annotation to pre-tokenized word-level units.

Initiatives such as the Indian Languages Corpora Initiative (ILCI) (Jha, 2012) and related efforts have contributed significantly to the development of annotated datasets and guidelines for Indian languages. However, these approaches are often dataset- or task-specific, and their reliance on token- or chunk-level annotation abstracts away character-level complexities inherent in Indic scripts.

When general-purpose annotation tools are applied to Indic data, several issues arise. Offset calculations based on bytes or code points can become inconsistent with the visually rendered text, leading to incorrect span boundaries, cursor misalignment during selection, and broken highlighting. These issues are closely related to Unicode text segmentation and grapheme cluster handling (The Unicode Consortium, 2022b,a). Similar challenges have also been observed in multilingual NLP settings where script-specific properties are not properly handled (Joshi et al., 2020).

Conjunct characters and ligatures may be incorrectly split, and tokenization methods designed for Latin scripts often produce inconsistent segmentation for Indic languages. This limitation has been noted in prior work on Indian language processing, where language-specific tokenization and normalization are required for accurate analysis (Kunchukuttan, 2020; Bharati et al., 1996). In addition, many annotation tools enforce predefined tokenization schemes that do not adapt well to morphologically rich languages (Tsarfaty et al., 2010), further reducing annotation flexibility.

Overall, existing tools and efforts do not provide a unified solution that simultaneously handles Unicode-aware rendering, grapheme-consistent offsets, and flexible annotation workflows for Indic languages. This gap makes it difficult to cre-

ate high-quality annotated corpora in such settings. Our work addresses these challenges by designing an annotation platform that aligns with the properties of Indic scripts. By using native Unicode rendering, maintaining a single canonical text, and avoiding forced tokenization, the system ensures accurate selection, stable offsets, and consistent annotation across multiple layers.

### 3 Challenges of Indic Script

Indic scripts are composed of grapheme clusters, where a single visible character may consist of multiple Unicode code points. Proper segmentation, therefore, requires grapheme-level processing rather than code-point-based indexing (The Unicode Consortium, 2022b,a; Ansary et al., 2024). Indic scripts pose inherent challenges that generic tools do not account for:

**Characters visually reorder:** Storage order of the graphemes may not be the same as the display order. Mostly, the vowel signs and combining marks are stored after the base consonant, but they can be rendered before, above, or below it. Hence, byte or code-point indices do not always correspond to the visually perceived character sequence. For example, in Bengali the grapheme ড়া is internally stored as consonant + nukta + vowel sign, while visually rendered as a single grapheme cluster with reordered components.

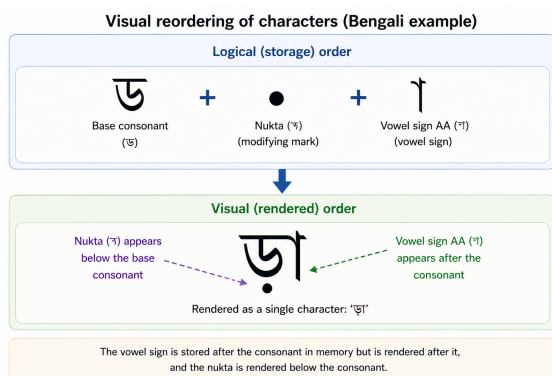


Figure 1: Example of visual reordering in Bengali grapheme rendering.

**Many glyphs are formed by multiple Unicode code points:** what appears as one “character” (one grapheme cluster) may be several code points (e.g. base consonant + nukta + vowel sign, such as Bengali ড়া = ড + ◌্ + া). Treating each code point as a unit breaks selection and alignment.

**Vowels attach in multiple positions:** Vowel signs can appear before, after, above, or below the con-

sonant, so a single “character” in the user’s mind corresponds to a variable-length sequence in memory.

**Segmentation of grapheme clusters is non-trivial:** Identifying user-perceived character boundaries requires Unicode grapheme cluster rules; naïve substring or tokenization often splits or merges them incorrectly.

An annotation tool for Indic languages must therefore: (1) use a rendering model that respects Unicode (e.g. native browser rendering); (2) store and interpret offsets in a single, stable string that matches what is displayed; (3) avoid forced tokenization that could split grapheme clusters or words incorrectly; and (4) support standard interchange formats like CoNLL-U (Nivre et al., 2016) while preserving correct character boundaries.

### 4 System Features

Our tool is designed so that Indic text is rendered, selected, and annotated without the issues described above. The following points summarize the main design choices.

**Native Unicode Rendering** The tool uses native browser Unicode rendering, which fully supports Indic scripts. It does not rely on a custom or simplified rendering engine; instead, it leverages the same rendering stack used by modern browsers for complex web scripts. As a result, issues related to reordering and ligatures that arise when display order is handled separately from storage are naturally mitigated.

**Accurate Selection and Highlighting** The tool provides accurate text selection and highlighting for Telugu, Hindi, Tamil, Malayalam, Bengali, Odia, Marathi, and Kannada across multiple annotation views (POS, span-rel, and SRL). After fetching doc.content, the frontend constructs a single client-side representation by collapsing consecutive whitespace and trimming leading/trailing spaces, and then derives sentence boundaries using a punctuation-based sentence segmentation strategy after normalization (with a punctuation-based fallback otherwise, including Indic sentence terminators). Annotation start/end offsets are computed against the same representation that is used to render sentence containers and highlight spans. As a result, the browser selection ranges align with the stored offsets used for highlighting, enabling stable and correct span rendering for Indic text in all supported layers.

## 5 Implementation & Features

Section 4 described the design principles that guide IndiAnn, particularly its Unicode-aware rendering and offset-preserving annotation model. In this section, we describe the core implementation choices that realize these principles in practice, the annotation layers currently supported by the platform, and the architectural flexibility that allows the framework to be extended to new annotation tasks while maintaining consistent alignment with the underlying text.

### 5.1 Core Implementation Design

**Offset Storage and Grapheme Clusters** The tool stores start/end offsets in a single offset space. All annotations (POS, spans, relations, dependencies) refer to this same string. During rendering, the client derives sentence containers and highlight spans using a whitespace-normalized representation of the stored content, so offset computations are performed against the same representation used for display. Offsets are computed in a manner that preserves grapheme clusters. Consequently, the substring boundaries used for highlighting match the mapped offset ranges, preventing script breakage due to unintended splitting or merging of grapheme clusters.

**No Forced Tokenization** The tool does not impose forced tokenization in the frontend or display pipeline. Instead of enforcing a fixed token grid that could split conjuncts or combine unrelated units, sentence boundaries are provided by the backend, but the frontend reconstructs sentence spans using a lightweight punctuation-based segmentation strategy. Within this structure, users can select spans corresponding to natural linguistic units (words, morphemes, or multi-word expressions), thereby avoiding the boundary errors typical of tokenizers designed for Latin scripts.

**CoNLL-U Import/Export and Character Boundaries** CoNLL-U import and export are supported while maintaining correct character boundaries. During import, the CoNLL-U file is converted into a canonical plain-text representation (e.g., one sentence per line), and token offsets are recomputed in this representation so that POS and dependency annotations align with the same underlying text. During export, annotations are merged back or generated from this canonical form. Consequently, CoNLL-U round-trips do not introduce offset drift or misaligned spans.

**Alignment of All Annotation Types** All annotation types (POS, spans, relations, dependencies) are stored using the same start/end offset model and align precisely with the underlying Indic text. A single offset space is maintained for the document, eliminating the possibility of inconsistencies that arise from separate token-index spaces. This unified representation ensures that POS tags, span labels, and dependency arcs remain correctly aligned with their corresponding character ranges.

**Backend:** FastAPI (Python) with PostgreSQL for persistence. Documents are stored with a single canonical content string; annotations store `start_offset` and `end_offset` into this string.

**Frontend:** Static HTML/JavaScript/CSS. The document content is fetched from the API and rendered in the browser. Before computing sentence containers and offsets for highlighting, the frontend collapses consecutive whitespace and trims leading/trailing spaces to ensure a consistent representation. Sentence boundaries are derived using a lightweight punctuation-based segmentation strategy, which includes Indic sentence terminators.

**Document Upload:** For plain text, the backend applies language-aware sentence tokenization (e.g. NLTK (Loper and Bird, 2002) for general languages, `indictoken` (Kunchukuttan, 2020) for Indic languages: Telugu, Hindi, Tamil, Malayalam, Bengali, Odia, Marathi, and Kannada). The result is stored as one sentence per line. For CoNLL-U, the file is parsed; a canonical plain-text form is derived; and token offsets are recomputed in that form so that all annotations refer to the same stored content.

Table 1 summarizes the design choices of IndiAnn.

### 5.2 Supported Annotation Tasks

All annotation layers are represented using a unified offset-based scheme over a single canonical text, where each annotation is stored with `start_offset` and `end_offset`, ensuring consistent alignment across layers.

*Part-of-speech (POS) Tagging:* annotation of tokens or spans with grammatical categories such as nouns, verbs, and adjectives, based on configurable tagsets.

*NER Tagging and Relation Annotation:*

- **Spans:** generic span annotations just like NER with configurable labels.

Choice	Description
Canonical text	All annotations are defined using character offsets with respect to a single, consistent document representation.
Content Immutability	The document text is not modified after upload; the same content is retrieved without any transformation.
Frontend whitespace normalization	The frontend collapses consecutive whitespace and trims leading/trailing spaces before rendering and computing sentence/offset positions, ensuring offsets remain consistent with the displayed representation.
Backend sentence segmentation	Sentence segmentation is performed during upload or import; the frontend derives sentence structure using a lightweight punctuation-based segmentation strategy over normalized content.
CoNLL-U Alignment	During CoNLL-U import, token offsets are recomputed within the canonical text to ensure alignment of POS and dependency annotations.

Table 1: Design choices of IndiAnn.

- **Relations:** links between spans or tokens, supporting relation extraction tasks.

*Dependency Relation Tagging:* annotation of syntactic head–dependent relations, where each token is linked to a governing word (head) and assigned a dependency label; compatible with CoNLL-U through token and head indices. This task is integrated within the POS layer (See Figure C.12 of Appendix C).

*Semantic Role Labelling (SRL):* annotation of predicate–argument structures, where arguments are associated with predicates and represented using character offsets.

All of these use the same document content and offset space, so they stay aligned with Indic text when the pipeline respects grapheme clusters and avoids forced tokenization.

### 5.3 Extensibility of Annotation Layers

IndiAnn follows a modular layer-oriented architecture in which each annotation layer is defined by a task-specific label schema, an offset-based representation over canonical text, and lightweight frontend interaction logic. Since all layers share the same `start_offset/end_offset` alignment model, new annotation layers can be introduced without modifying the core rendering, storage, or curation pipeline.

Beyond the currently supported layers the framework can be extended to additional span-level classification, event annotation, discourse relation labeling, coreference annotation, and other structured linguistic tasks through project-level label configuration and minor interface extensions.

At present, IndiAnn is best suited for text-centric annotation tasks grounded in character-offset spans or relations over canonical text. Tasks requiring multimodal alignment, deeply nested

graph structures, or specialized preprocessing may require additional modeling beyond the current implementation.

### 5.4 Sentence Segmentation and Scalability

Sentence segmentation is performed once at upload time. For plain text documents, the backend applies language-aware sentence tokenization (e.g., NLTK for non-Indic languages and the Indic NLP Library for Indic languages), then joins the resulting sentences using newline characters to form a canonical string. For CoNLL-U documents, each sentence block is converted into a canonical plain-text form by representing it as a sequence of whitespace-separated tokens. Token offsets are then recomputed over this representation to ensure alignment with the underlying text.

Upon retrieval, the API returns this canonical content; when it is already sentence-separated, the frontend derives sentence containers using punctuation-based segmentation over the normalized text. If newline segmentation is not available, the frontend uses a lightweight punctuation-based heuristic to derive sentence spans for rendering and offset mapping.

Because sentence boundaries are encoded as lightweight newline delimiters and annotation offsets are defined over a single immutable string, the system scales linearly with document length and annotation count. Rendering and interaction primarily involve substring operations and span insertions, which maintain responsiveness even for relatively large documents.

### 5.5 Sparse and Incremental Annotation

The tool supports sparse and incremental annotation. Annotators are not required to label every token or span in a document; instead, they can focus

on selected regions (e.g., complex constructions or specific phenomena) and extend coverage in subsequent passes. All annotation layers (POS, spans, relations, dependencies, SRL) are stored independently as records keyed by document ID and character offsets into the canonical text. As a result, additional annotations can be introduced without modifying or recomputing existing ones, enabling efficient incremental annotation workflows while preserving offset consistency.

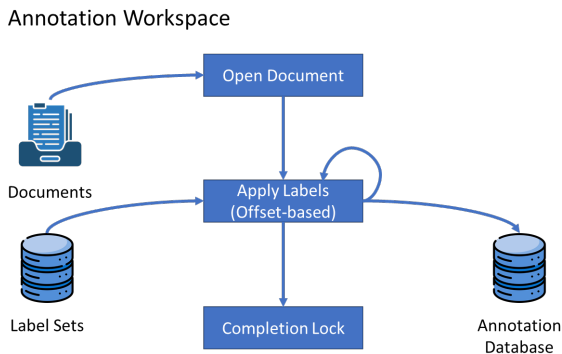


Figure 2: Workflow Diagram for Annotator

## 5.6 Three-Tier Management

The platform supports a collaborative annotation workflow with distinct roles. The graphical user interface can be described as three role-oriented views: (i) an annotation view for annotators, (ii) a curation view for curators to resolve conflicts and consolidate decisions, and (iii) an administration view to configure projects, labels, and membership.

**Annotator** Annotators perform primary annotation on assigned documents. In the annotation view, the user is presented with the document text rendered using native Unicode support and a compact control panel for labeling actions. Annotation is performed directly on the text: the annotator selects a span and assigns a label; the tool stores start/end offsets relative to the canonical document content, ensuring that selections remain stable for Indic grapheme clusters. The annotator’s workflow is visualized in the Figure 2.

Typical annotator actions include:

- **Labeling:** Apply labels by selecting text and choosing a label from the available tagset.
- **Deletion and correction:** Delete incorrect annotations (including, where applicable, associated dependent structures such as SRL

roles attached to a deleted predicate) and re-annotate the corrected span. The system does not currently support direct in-place editing of annotations; corrections require manual re-selection of the span (Figure 3).

- **Multi-layer annotation:** Work across multiple layers (POS, spans/entities, relations, dependencies, SRL) over the same canonical text, depending on the project configuration. It is displayed in Figure 4.
- **Export:** After annotations are done, the annotated file can be exported using standard formats like JSON (See Figure B.5 of Appendix B) and CoNLL-U (See Figure B.6 of Appendix B) for dependency/POS workflows where applicable.
- **Progress control:** There is a provision of progress control to mark documents as complete to reduce accidental edits.

**Curator** Curators validate, resolve conflicts, and build a higher-quality consensus layer over annotator submissions. The curation view aggregates all annotations that share the same *(start, end)* span and summarizes them *by label*. For example, the same highlighted word or span might have been annotated as *noun* by five annotators and as *verb* by three annotators. For each competing label, the interface provides explicit curation actions (approve/reject), allowing curators to record a decision at the level of label groups rather than inspecting annotations one by one. Curation itself is only made available once all assigned annotators have *locked* (marked as completed) their work for a document; this gating is configured and monitored by the admin role. The curation workflow is visualized in the Figure 5. This figure shows a curation tab of an NER layer (Figure 6)

Typical curator actions include:

- **Conflict inspection:** view all labels proposed by different annotators for the same *(start, end)* span.
- **By-label breakdown:** inspect per-label vote counts for a span (e.g., *noun: 5, verb: 3*), making disagreement patterns immediately visible.
- **Approve/reject per label:** for each competing label group (e.g., *noun* and *verb*), use dedicated **Approve** and **Reject** buttons to accept

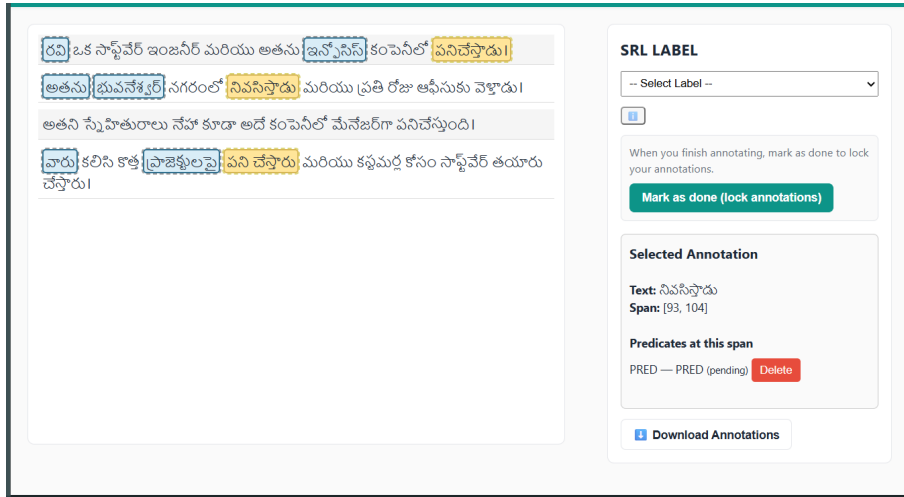


Figure 3: SRL Annotation tab of an Annotator. The document is written in Telugu. The orange coloured box is a predicate while the blue boxes are roles.

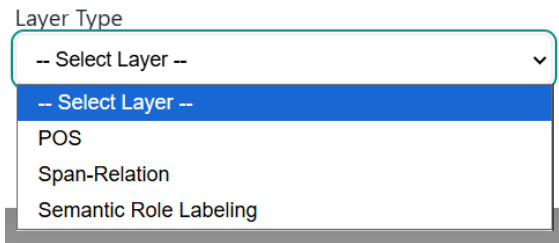


Figure 4: A Snapshot showing provision of multi-layer annotation

or discard that label for the current span; this yields a curated consensus decision.

- **Quality control:** identify systematic problems (e.g. recurrent boundary mistakes) and correct them while keeping the canonical text fixed.
- **Error-focused refinement:** operationalize script-induced issues (offset/boundary errors) as explicit curation decisions.

**Admin** Admins configure projects and manage users and resources needed to run annotation campaigns. The administration view supports creating projects, defining label sets, assigning user roles, and managing documents and exports. Figure A.1 (Appendix A) shows the admin dashboard of **Indi-Ann**. Typical admin actions include:

- **Project setup:** create projects and configure which annotation layers are enabled (e.g. POS vs. Span-Relation vs. SRL). See Figure 4.
- **Label management:** define the permitted labels/tagsets per project and update them as guidelines evolve. See Figure A.2 of Appendix A.

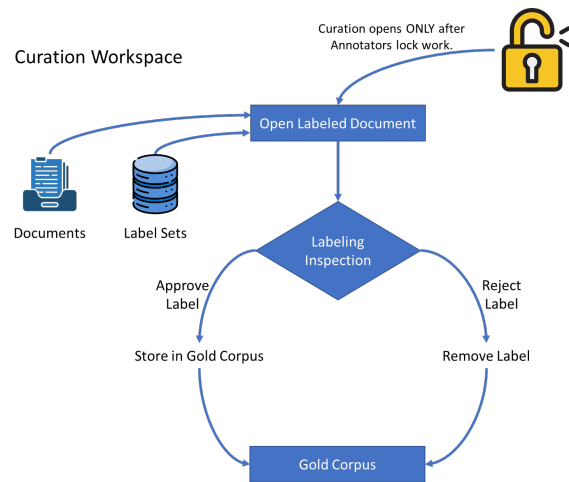


Figure 5: Workflow Diagram for Curator

- **User and role management:** add users to projects and assign roles (annotator/curator/admin). See Figure A.3 of Appendix A.
- **Document management:** upload documents (Plain Text/ CoNLL-U/Word document/PDF), monitor progress, and manage exports.
- **Completion locks:** monitor whether annotators have locked (completed) a document, and if needed, *reverse/unlock* a lock to allow further edits (e.g., after guideline updates or detected errors). See Figure A.4 of Appendix A.
- **Curation gating:** enforce or configure the policy that curation starts only once all assigned annotators have locked their work, so

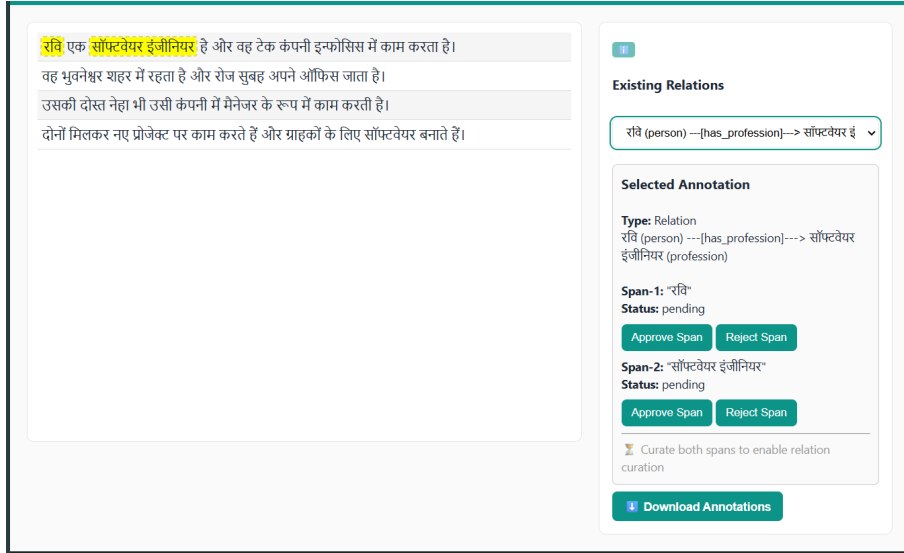


Figure 6: NER Curation tab of a Curator. The document is written in Hindi Language.

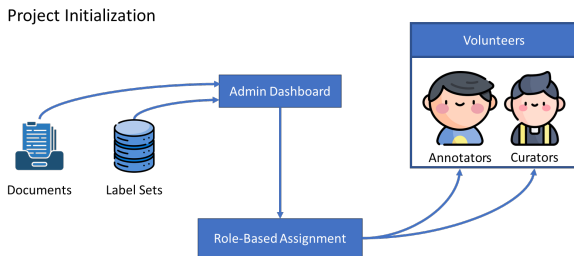


Figure 7: Workflow Diagram for Admin

curators can curate a stable snapshot of annotations.

- **Governance:** enforce workflow policies appropriate for a dataset (e.g. completion locks, curation requirements) and maintain traceability of campaign progress.

The typical workflow for admins is presented in Figure 7.

## 6 Linguistic Coverage

The tool is designed for low-resource Indic languages, including Telugu, Hindi, Tamil, Malayalam, Bengali, Odia, Marathi, and Kannada. More generally, it supports scripts that follow standard Unicode encoding and use explicit whitespace for word delimitation. This design makes the platform applicable to most Indic scripts used in Indo-Aryan and Dravidian languages, in which grapheme clusters align with user-perceived character units and can be handled reliably via native rendering and offset-based annotation. However, linguistic phenomena that obscure word boundaries (e.g., sandhi or compounding) may require additional handling beyond the current design.

Although IndiAnn was developed primarily for Indic languages, its design is not restricted to Indic scripts. By relying on native Unicode rendering, grapheme-consistent offset storage, and direct span selection over canonical text, the framework can support other complex scripts as well. In preliminary experiments, we observed stable rendering and annotation behavior for Arabic and Urdu, suggesting that extending IndiAnn to additional Unicode-compliant scripts would require minimal changes beyond language-specific preprocessing such as sentence segmentation.

## 7 Error Mitigation

IndiAnn is designed to reduce annotation errors that arise when software assumptions about text representation do not align with the properties of Unicode-complex scripts. By relying on native Unicode rendering, preserving annotations in a single canonical offset space, and avoiding forced tokenization in the annotation interface, the framework minimizes cursor-offset mismatch, broken span boundaries, and annotation inconsistencies caused by grapheme fragmentation

## 8 Conclusion

We have proposed an annotation platform (**IndiAnn**), built for low-resource Indic languages. The key features of this framework are: (1) reliance on native Unicode rendering and a single canonical text with offset-based annotations that preserve grapheme clusters; (2) no forced tokenization in the user interface, avoiding script breakage. Together, these properties allow accu-

rate selection, highlighting, and export for Telugu, Hindi, Tamil, Malayalam, Bengali, Odia, Marathi, and Kannada, addressing the limitations of tools that are built mainly for European scripts and that suffer from vowel matra reordering, conjunct handling, multi-byte Unicode, and tokenization errors. **IndiAnn** supports annotation for four different tasks such as POS tagging, NER, Dependency Relation extraction and Semantic Role Labeling. This framework would be the stepping stone towards seamless annotation experience for low-resource Indic Languages.

## Limitations

While the proposed platform is designed to support annotation for low-resource Indic languages, our framework has a few limitations. First, the proposed framework supports multiple annotation layers, but they are handled within a shared workspace, which may increase interface complexity and cognitive load in multi-layer projects. Second, the framework assumes a stable preprocessing pipeline for sentence segmentation and offset alignment. Changes in preprocessing may affect consistency and require careful version control. Third, while the platform handles Unicode-level complexities such as grapheme clusters and rendering across Indic scripts, it does not explicitly address linguistic phenomena such as sandhi and compounding (samās) yet, where multiple lexical units are realized as a single orthographic word. In such cases, boundaries required for annotation may not align with whitespace-delimited tokens, potentially limiting fine-grained annotation. Handling such cases is our immediate future work.

## References

- Nazmuddoha Ansary, Quazi Adibur Rahman Adib, Tahsin Reasat, Asif Shahriyar Sushmit, Ahmed Intiaz Humayun, Sazia Mehnaz, Kanij Fatema, Mohammad Mamun Or Rashid, and Farig Sadeque. 2024. Unicode normalization and grapheme parsing of indic languages. In *Proceedings of the 2024 joint international conference on computational linguistics, language resources and evaluation (LREC-COLING 2024)*, pages 17019–17030.
- Rafiya Begum, Samar Husain, Arun Dhvaj, Dipti Misra Sharma, Lakshmi Bai, and Rajeev Sangal. 2008. Dependency annotation scheme for indian languages. In *Proceedings of the Third International Joint Conference on Natural Language Processing: Volume-II*.
- Emily M Bender. 2011. On achieving and evaluating language-independence in nlp. *Linguistic Issues in Language Technology*, 6.
- Akshar Bharati, Vineet Chaitanya, and Rajeev Sangal. 1996. *Natural Language Processing: A Paninian Perspective*. Prentice-Hall of India.
- Steven Bird, Ewan Klein, and Edward Loper. 2009. *Natural language processing with Python: analyzing text with the natural language toolkit*. ”O’Reilly Media, Inc.”.
- Girish Nath Jha. 2012. The tdil program and the indian language corpora initiative. In *Language resources and evaluation conference*.
- Pratik Joshi, Sebastin Santy, Amar Budhiraja, Kalika Bali, and Monojit Choudhury. 2020. The state and fate of linguistic diversity and inclusion in the nlp world. In *Proceedings of the 58th annual meeting of the association for computational linguistics*, pages 6282–6293.
- Jan-Christoph Klie, Michael Bugert, Beto Boullosa, Richard Eckart De Castilho, and Iryna Gurevych. 2018. The inception platform: Machine-assisted and knowledge-oriented interactive annotation. In *Proceedings of the 27th international conference on computational linguistics: system demonstrations*, pages 5–9.
- Ritesh Kumar, Shiv Bhusan Kaushik, Pinkey Nainwani, and Girish Nath Jha. 2021. [Creating and managing a large annotated parallel corpora of indian languages](#). *Preprint*, arXiv:2112.01764.
- Anoop Kunchukuttan. 2020. The IndicNLP Library. [https://github.com/anoopkunchukuttan/indic\\_nlp\\_library/blob/master/docs/indicnlp.pdf](https://github.com/anoopkunchukuttan/indic_nlp_library/blob/master/docs/indicnlp.pdf).
- Edward Loper and Steven Bird. 2002. Nltk: The natural language toolkit. In *Proceedings of the ACL-02 Workshop on Effective tools and methodologies for teaching natural language processing and computational linguistics*, pages 63–70.
- Arijit Maji, Raghvendra Kumar, Akash Ghosh, Nemil Shah, Abhilekh Borah, Vanshika Shah, Nishant Mishra, Sriparna Saha, and 1 others. 2025. Drishtikon: A multimodal multilingual benchmark for testing language models’ understanding on indian culture. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 1289–1313.
- Hiroki Nakayama, Takahiro Kubo, Junya Kamura, Yasufumi Taniguchi, and Xu Liang. 2018. [doccano: Text annotation tool for human](#). Software available from <https://github.com/doccano/doccano>.
- Joakim Nivre, Marie-Catherine De Marneffe, Filip Ginter, Yoav Goldberg, Jan Hajic, Christopher D Manning, Ryan McDonald, Slav Petrov, Sampo Pyysalo, Natalia Silveira, and 1 others. 2016. Universal

- dependencies v1: A multilingual treebank collection. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC'16)*, pages 1659–1666.
- Tjong Kim Sang and De Meulder. 2003. Introduction to the conll-2003 shared task. In *CoNLL*.
- Pontus Stenetorp, Sampo Pyysalo, Goran Topić, Tomoko Ohta, Sophia Ananiadou, and Jun'ichi Tsujii. 2012. Brat: a web-based tool for nlp-assisted text annotation. In *Proceedings of the Demonstrations at the 13th Conference of the European Chapter of the Association for Computational Linguistics*, pages 102–107.
- The Unicode Consortium. 2022a. [Unicode standard annex #29: Unicode text segmentation, version 15.0](#).
- The Unicode Consortium. 2022b. *The Unicode Standard, Version 15.0*. The Unicode Consortium, Mountain View, CA.
- Maxim Tkachenko, Mikhail Malyuk, Andrey Holmanyuk, and Nikolai Liubimov. 2020-2025. [Label Studio: Data labeling software](#). Open source software available from <https://github.com/HumanSignal/label-studio>.
- Reut Tsarfaty, Djamé Seddah, Yoav Goldberg, Sandra Kübler, Yannick Versley, Marie Candito, Jennifer Foster, Ines Rehbein, and Lamia Tounsi. 2010. Statistical parsing of morphologically rich languages (spmrl) what, how and whither. In *Proceedings of the NAACL HLT 2010 First Workshop on Statistical Parsing of Morphologically-Rich Languages*, pages 1–12.
- Seid Muhie Yimam, Iryna Gurevych, Richard Eckart De Castilho, and Chris Biemann. 2013. Webanno: A flexible, web-based and visually supported system for distributed annotations. In *Proceedings of the 51st annual meeting of the Association for Computational Linguistics: system demonstrations*, pages 1–6.

# Appendix

## A Provision of Admin

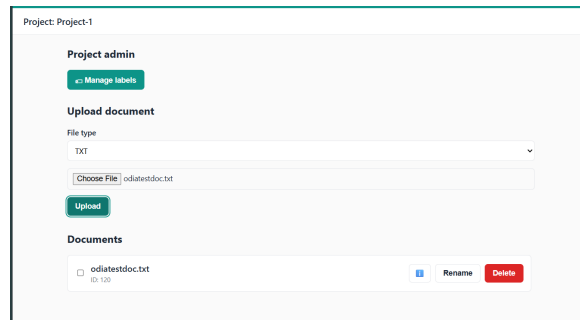


Figure A.1: A snapshot of Admin Dashboard

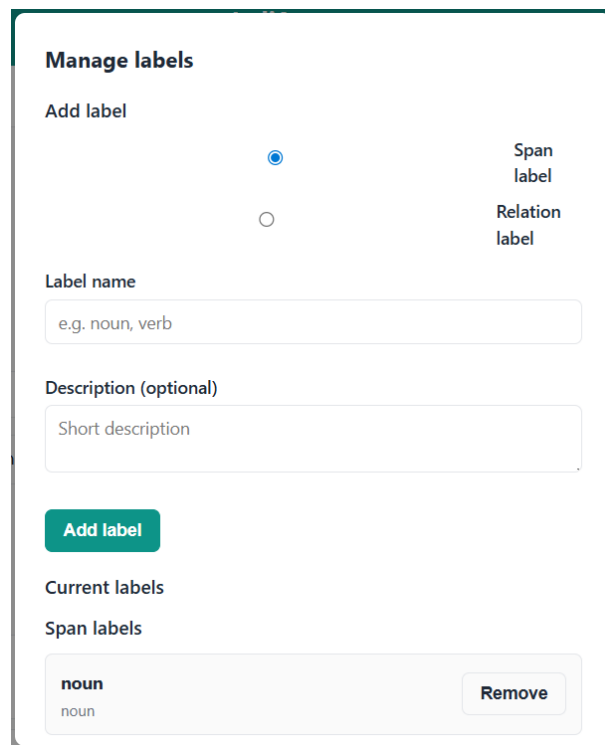


Figure A.2: A snapshot of label management provision

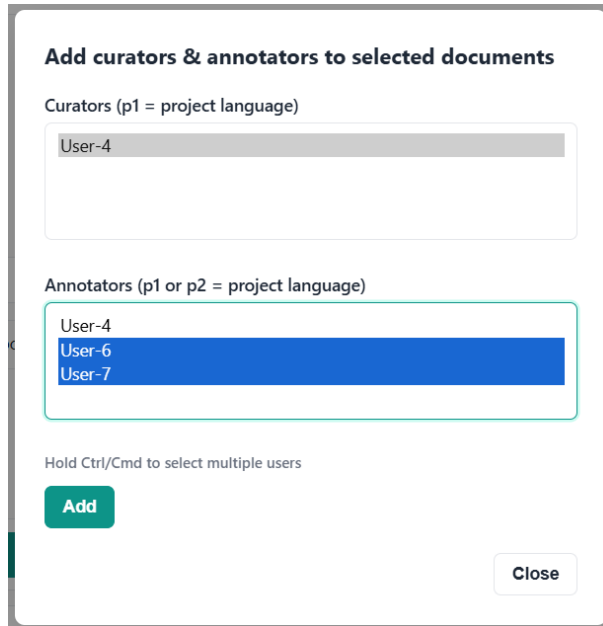


Figure A.3: A snapshot of role management provision

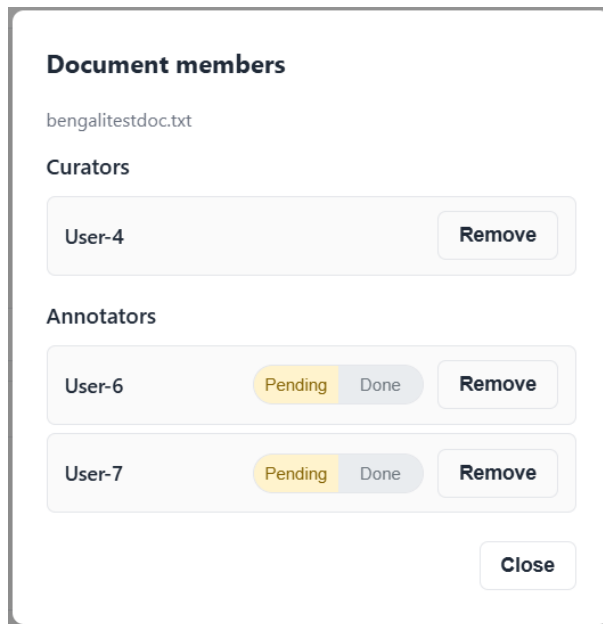


Figure A.4: A snapshot of lock management provision

## B File Export Format

```
{
  "document_id": 10,
  "project_id": 4,
  "filename": "telugutestdoc.txt",
  "text": "రవి ఒక సాఫ్ట్‌వేర్ ఇంజనీర్ మరియు అతను ఇన్ఫోసీస్ కంపెనీలో పనిచేస్తాడు।\nఅతను భువనేశ్వర్ తయారు చేస్తారు।",
  "predicates": [
    {
      "id": 3,
      "predicate_label": "PRED",
      "predicate_text": "పనిచేస్తాడు।",
      "start_offset": 57,
      "end_offset": 69,
      "roles": [
        {
          "id": 5,
          "role_label": "arg0",
          "role_text": "రవి",
          "start_offset": 0,
          "end_offset": 3
        },
        {
          "id": 6,
          "role_label": "argm-loc",
          "role_text": "ఇన్ఫోసీస్",
          "start_offset": 38,
          "end_offset": 47
        }
      ]
    }
  ]
}
```

Figure B.5: Exported JSON file of an SRL-annotated document. The document is written in Telugu language.

1	রাহুল	রাহুল	pronoun	pronoun	-	-	-	-	start=0 end=5
2	একজন	একজন	-	-	-	-	-	-	start=6 end=10
3	সফটওয়্যার	সফটওয়্যার	noun	noun	-	-	-	-	start=11 end=21
4	ইঞ্জিনিয়ার	ইঞ্জিনিয়ার	noun	noun	-	-	-	-	start=22 end=33
5	এবং	এবং	conjunction	conjunction	-	-	-	-	start=34 end=37
6	সে	সে	pronoun	pronoun	-	-	-	-	start=38 end=40
7	একটি	একটি	-	-	-	-	-	-	start=41 end=45
8	বড়	বড়	-	-	-	-	-	-	start=46 end=49
9	কোম্পানিতে	কোম্পানিতে	-	-	-	-	-	-	start=50 end=60
10	কাজ	কাজ	verb	verb	-	-	-	-	start=61 end=64
11	করে।	করে।	verb	verb	-	-	-	-	start=65 end=69
12	সে	সে	pronoun	pronoun	-	-	-	-	start=70 end=72
13	কলকাতা	কলকাতা	-	-	-	-	-	-	start=73 end=79
14	শহরে	শহরে	-	-	-	-	-	-	start=80 end=84
15	থাকে	থাকে	-	-	-	-	-	-	start=85 end=89
16	এবং	এবং	-	-	-	-	-	-	start=90 end=93
17	প্রতিদিন	প্রতিদিন	adverb	adverb	-	-	-	-	start=94 end=102
18	অফিসে	অফিসে	noun	noun	-	-	-	-	start=103 end=108
19	যায়।	যায়।	-	-	-	-	-	-	start=109 end=114
20	তার	তার	-	-	-	-	-	-	start=115 end=118
21	বন্ধু	বন্ধু	-	-	-	-	-	-	start=119 end=124
22	অনিতা	অনিতা	-	-	-	-	-	-	start=125 end=130
23	একই	একই	-	-	-	-	-	-	start=131 end=134
24	কোম্পানিতে	কোম্পানিতে	-	-	-	-	-	-	start=135 end=145
25	ম্যানেজার	ম্যানেজার	-	-	-	-	-	-	start=146 end=155
26	হিসেবে	হিসেবে	-	-	-	-	-	-	start=156 end=162
27	কাজ	কাজ	-	-	-	-	-	-	start=163 end=166
28	করে।	করে।	-	-	-	-	-	-	start=167 end=171
29	তারা	তারা	-	-	-	-	-	-	start=172 end=176
30	একসাথে	একসাথে	-	-	-	-	-	-	start=177 end=183
31	নতুন	নতুন	-	-	-	-	-	-	start=184 end=188
32	প্রজেক্টে	প্রজেক্টে	-	-	-	-	-	-	start=189 end=198
33	কাজ	কাজ	-	-	-	-	-	-	start=199 end=202
34	করে	করে	-	-	-	-	-	-	start=203 end=206
35	এবং	এবং	-	-	-	-	-	-	start=207 end=210
36	গ্রাহকদের	গ্রাহকদের	noun	noun	-	-	-	-	start=211 end=220
37	জন্য	জন্য	-	-	-	-	-	-	start=221 end=225
38	সফটওয়্যার	সফটওয়্যার	-	-	-	-	-	-	start=226 end=236
39	তৈরি	তৈরি	verb	verb	-	-	-	-	start=237 end=241
40	করে।	করে।	verb	verb	-	-	-	-	start=242 end=246

Figure B.6: Exported CoNLL-U file of a POS annotated document. The document is written in Bengali language.

## C Sample Snapshots of IndiAnn

This section contains images of documents annotated using POS tags for the supported languages.



Figure C.7: POS Annotated Odia document



Figure C.8: POS Annotated Tamil document



Figure C.9: POS Annotated Bengali document

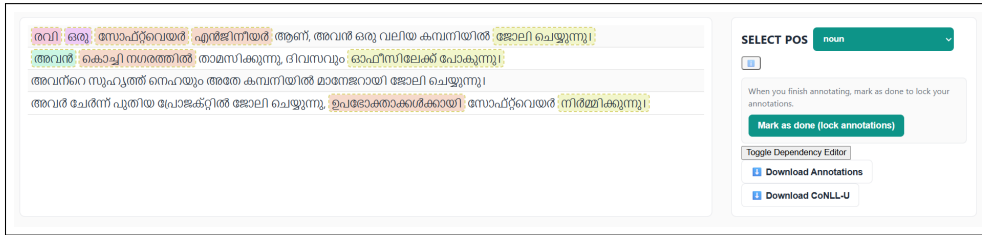


Figure C.10: POS Annotated Malayalam document

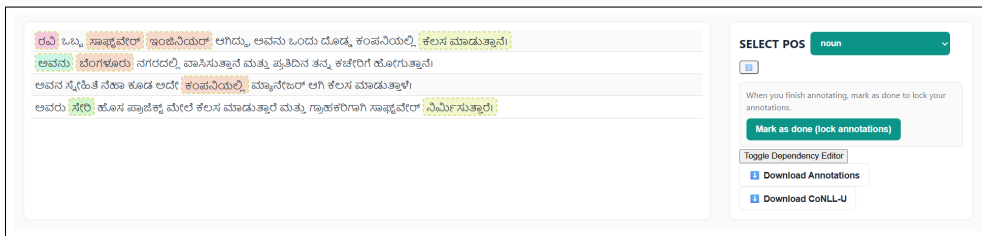


Figure C.11: POS Annotated Kannada document

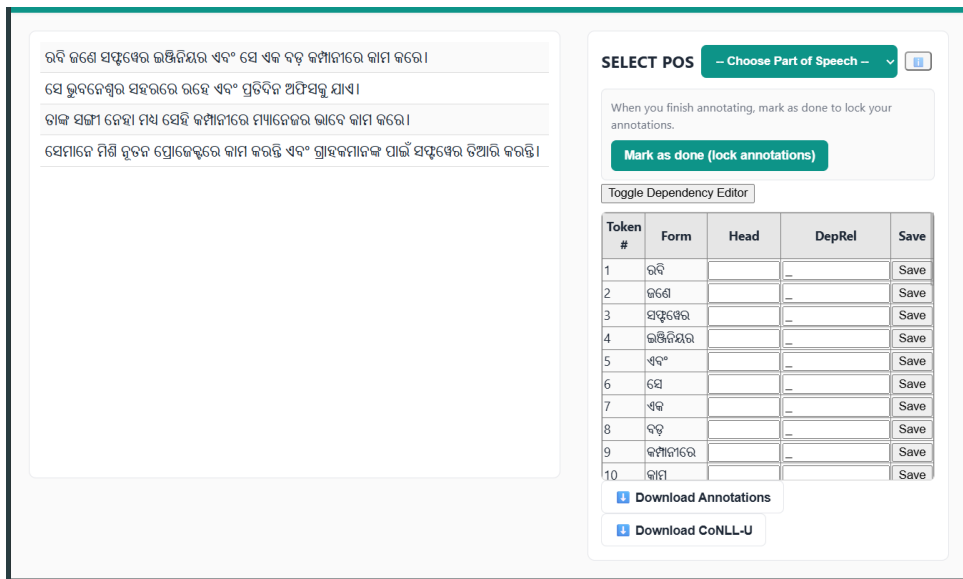


Figure C.12: Dependency tagging of an odia Document

# Designing Annotation Guidelines for Trait-Based Arabic Automated Essay Scoring: A Systematic Methodology

Walid Massoud<sup>1</sup> Houda Bouamor<sup>2</sup>  
Abdelrahman Abdel Latif Hussein<sup>3</sup> Abdullah Mohamed Mohamed Zekri<sup>4</sup>

<sup>1</sup>Qatar University

<sup>2</sup>Carnegie Mellon University in Qatar

<sup>3</sup>Ministry of Education, Egypt

<sup>4</sup>National Center for Examinations and Educational Evaluation, Egypt

## Abstract

Automated Essay Scoring (AES) fundamentally depends on high-quality annotated data, yet systematic approaches to developing annotation guidelines remain largely undocumented, especially for Arabic. We present a comprehensive methodology for trait-based Arabic AES annotation, applied to build a dataset of 7,859 essays by high school students annotated across seven writing traits, achieving substantial inter-annotator agreement (QWK: 0.66–0.75). Our methodology encompasses: (1) a seven-dimensional scoring framework grounded in Arabic linguistic and rhetorical conventions; (2) over 25 pages of Arabic-language guidelines with terminology unification, text-type-specific scoring descriptors, and annotated student examples; (3) a multi-stage training protocol that raised annotator agreement before production began; and (4) quality assurance mechanisms, including dual annotation and supervisor adjudication. We release all materials publicly, providing both a validated foundation for Arabic AES research and a replicable template for annotation guideline development in other morphologically complex, under-resourced languages

## 1 Introduction

Automated Essay Scoring (AES) has emerged as a critical NLP application enabling scalable writing assessment. While significant progress has been made for English (Mathias and Bhattacharyya, 2018; Crossley et al., 2023b), its efficiency fundamentally depends on the quality of the human-annotated data used for model training. In practice, this quality is shaped by the extent to which human evaluators achieve objectivity, accuracy, and inter-rater reliability. Essay assessment, by its nature, involves subjective judgment and interpretation, making it inherently difficult to ensure consistency and precision across raters. These challenges are further amplified in multi-trait scoring settings, where

evaluators must simultaneously assess multiple dimensions of writing quality, increasing cognitive load and the potential for inconsistency. Ensuring the reliability and validity of the resulting scores remains a persistent challenge. This problem is especially pronounced for Arabic, an under-resourced language with distinct morphological, syntactic, and rhetorical characteristics that complicate both writing assessment and annotation guideline design. In such contexts, annotation is not merely a labeling task, but a structured reasoning process that must be explicitly supported to achieve consistent and reliable outcomes.

Arabic AES research faces particular challenges: scarcity of annotated datasets, the lack of established annotation protocols (Bashendy et al., 2024), and the inherent linguistic complexity of Modern Standard Arabic (MSA). MSA exhibits complex morphological agreement, flexible word order, and a rich system of cohesive devices that do not map onto frameworks developed for European languages. Existing Arabic AES datasets provide limited annotation documentation (Habash and Palfreyman, 2022), making replication and adaptation difficult. Annotators must distinguish subtle gradations - such as the difference between surface Arabic cohesive devices (روابط خطية) and deep lexical cohesion (ترابط معجمي), while maintaining consistency across thousands of essays. Without systematic, Arabic-specific guidelines, this leads to low inter-annotator agreement and unreliable training data.

This paper presents a comprehensive methodology for developing annotation guidelines for trait-based Arabic AES, applied to build a large-scale Arabic essay dataset comprising 7,859 essays from 4,372 highschool students across 24 schools in an Arab country (Bashendy et al., 2025b). While the individual components of our approach (multi-stage annotator training, dual annotation, calibra-

tion sessions) reflect established best practices in annotation methodology (Artstein and Poesio, 2008; Williamson et al., 2012), their systematic adaptation to Arabic presents non-trivial challenges. Arabic-specific morphology, cohesive conventions, and rhetorical norms require decisions that cannot be derived from existing frameworks developed for English or other languages. Our core contribution is therefore the principled, documented *operationalization* of these practices for Arabic academic writing. Concretely, we contribute:

1. A seven-dimensional scoring framework grounded in Arabic linguistic and rhetorical conventions, with observable, Arabic-specific indicators replacing vague criteria (e.g., topic sentences *الجملة الموضوعية* as a proxy for organization);
2. Over 25 pages of annotation guidelines written entirely in Arabic, including a terminology unification section, text-type-specific scoring descriptors, and annotated student writing examples at each score level, none of which exist for Arabic AES in prior work;
3. A multi-stage training protocol with calibration providing evidence for each training decision rather than reporting protocol as a black box;
4. Quality assurance mechanisms adapted to the Arabic context, including adjudication protocols for code-switching and AI-generated text specific to MSA student writing; and
5. Full public release of all materials to enable direct replication and adaptation.

## 2 Related Work

**AES Annotation Frameworks.** English AES benefits from well-established datasets such as ASAP (Mathias and Bhattacharyya, 2018), ELIPSE (Crossley et al., 2023a), and PERSUADE (Crossley et al., 2023b), but these provide limited documentation of annotation development. Cross-lingual frameworks such as TCFLE-8 for French (Wilkens et al., 2023), MERLIN for European languages (Boyd et al., 2014), and ACEA for Chinese (He et al., 2022) offer holistic assessments but lack trait-level granularity. None addresses Arabic academic writing, its morphological richness, or Arabic rhetorical conventions.

**Annotation Guideline Development.** Best practices for subjective NLP annotation emphasize clear definitions, concrete examples, and iterative refinement (Artstein and Poesio, 2008). Detailed rubrics, borderline-case training, and multiple pilot rounds are essential (Fort et al., 2011). In educational assessment, high inter-rater reliability requires extensive training and calibration (Williamson et al., 2012), with feedback and group discussion significantly improving agreement (Landis and Koch, 1977).

**Arabic AES and Annotation.** Arabic AES has been constrained by limited annotated data. ZAE-BUC (Habash and Palfreyman, 2022) provides linguistic annotations but lacks trait-specific labels. QAES (Bashendy et al., 2024) introduced trait annotations at small scale (195 essays). The TAQEEM shared task (Bashendy et al., 2025a) focuses on dataset release rather than annotation methodology. To the best of our knowledge, no prior work has published Arabic annotation guidelines accounting for the full range of Arabic-specific writing features.

## 3 Annotation Framework Design

We designed a scoring framework that reflects the linguistic and rhetorical characteristics of Arabic academic writing, rather than adapting existing frameworks developed for other languages.

### 3.1 Design Principles

Our framework was guided by six principles: (1) **trait-based assessment**; (2) **text-type differentiation** between expository (*نص تفسيري*) and persuasive (*نص إقناعي*) writing; (3) **clear Arabic operationalization**, replacing vague criteria like “good organization” with observable indicators such as topic sentences (*الجملة الموضوعية*); (4) **granular but manageable scales**; (5) **evidence-based scoring** requiring textual justification; and (6) **cultural and linguistic appropriateness** for MSA norms and Arabic cohesive mechanisms, including reference (*الإحالة*), ellipsis/deletion (*الحذف*), substitution (*الإبدال*), and lexical cohesion (*الترابط المعجمي*).

### 3.2 Seven-Dimensional Scoring Framework

Table 1 summarizes our seven traits, adopted from the Core Academic Skills Test rubric (Bashendy

Trait	Scale	Description
REL	0–2	Relevance to Arabic prompt topic
ORG	0–5	Structure per Arabic essay conventions
VOC	0–5	Arabic lexical range, precision, MSA use
STY	0–5	Arabic cohesive devices & discourse patterns
DEV	0–5	Idea clarity, evidence, argument quality
MEC	0–5	Arabic spelling (إملاء), punctuation (ترقيم)
GRA	0–5	Syntactic variety & Arabic grammatical accuracy
HOL	0–32	Sum of all trait scores

Table 1: Seven-dimensional scoring framework. REL uses a 3-point scale; all other traits use 6-point scales (0–5). HOL is the sum of all traits.

et al., 2025b) and adapted for Arabic academic writing at the high-school level (grades 10–12).

**VOC** distinguishes limited vocabulary (المدى المحدود) from rich MSA usage, penalizing three error types: semantic mismatch (اختيار الكلمة غير المناسبة دلاليًا), collocation errors (أخطاء التراكيب المعجمية), and filler words (الكلمات الحشوية), while rewarding idiomatic expressions (التعبيرات الاصطلاحية) and implicit meaning (المعنى الضمني).

**STY** evaluates Arabic linear cohesion (الروابط الخطية) via four mechanisms: referencing (الإحالة), connective tools (الأدوات), deletion (الحذف), and substitution (الإبدال); and lexical cohesion (الترايب المعجمي) through repetition, synonymy, antonymy, and semantic field. The guidelines enumerate six Arabic organizational patterns (أنماط عرض الأفكار): cause-effect, compare-contrast, classification, chronological, interpretation, and pros-cons.

**MEC** covers Arabic-specific orthographic errors: hamza confusion (همزة القطع وألف الوصل), *ta marbuta* confusion (الناء المربوطة والهاء), and *ya* vs. *alif maqsura* (الياء والألف المقصورة).

**GRA** distinguishes simple Arabic structures (التراكيب البسيطة) from syntactic

variety (التنوع التركيبي) including conditional, relative, and parenthetical constructions. If REL = 0, all other traits are automatically scored 0.

### 3.3 Expository vs. Persuasive Differentiation

The most critical design decision was distinguishing evaluation criteria for the two text types. Arabic expository writing demands neutrality and objectivity (حياد وموضوعية), while persuasive writing requires a clear position (موقف واضح) supported by evidence and rhetorical technique.

For **DEV**, expository essays are evaluated on clarity of explanation, depth of analysis, and objectivity (الحياد التفسيري). The guidelines warn annotators that first-person phrases “I believe” (أنا أرى), and “we should” (يجب علينا) constitute a deficiency in expository writing. Persuasive essays are evaluated on consistency of position, argument strength, evidence quality, persuasive techniques (الأساليب الإقناعية), and acknowledgment plus refutation of counterarguments (عرض الآراء المختلفة).

## 4 Guideline Development

Translating the scoring framework into a usable annotation instrument required an iterative development process grounded in authentic Arabic student writing, expert consensus, and empirical piloting.

### 4.1 Development Process

Guideline development proceeded through six phases over four months: (1) Rubric adoption (Weeks 1 - 2); (2) Arabic-specific adaptation for Style and Development (Weeks 3 - 4); (3) Expert exemplar development with two Arabic educators independently annotating 60 essays to consensus (Weeks 5 - 8); (4) Guidebook drafting in Arabic (Weeks 9 - 12); (5) External blind review by an Arabic language pedagogy expert, who evaluated the guidelines and provided structured evaluation leading to iterative revisions (Weeks 13 - 14); and (6) Pilot testing and refinement with three annotators on a subset of 20 essays (Weeks 15 - 16).

### 4.2 Guideline Components

The final guidebook, written entirely in Arabic to eliminate translation ambiguity and to ensure that annotators engage with scoring criteria in the same language as the essays they evaluate, comprises

four components: Conceptual Foundations, Scoring Descriptors, Annotated Arabic Examples, and Edge Case Protocols.

**Conceptual Foundations** The guidebook opens with a terminology unification section (توحيد المفاهيم والمصطلحات). **Text type definitions** provide full Arabic characterizations of expository and persuasive writing with a comparison table and multiple Arabic student writing examples. **Vocabulary terminology** (المفردات) defines all error types, idiomatic expressions, and implicit meaning, each illustrated with authentic student Arabic. **Style and cohesion terminology** (الأسلوب والتماسك البنائي) defines all four linear cohesion mechanisms and all six organizational patterns with Arabic examples and characteristic connective words. **Grammar terminology** (البناء والتراكيب) distinguishes simple from complex Arabic structures.

**Scoring Descriptors** For each trait and score level, we provide detailed Arabic descriptors derived from the original rubrics, organized into three complementary components: descriptor, performance characteristics, and illustrative example. As shown in Table 2, this structure specifies observable and measurable criteria while grounding them in representative examples of expected student responses, enabling annotators to anchor their judgments in concrete textual evidence. In addition, targeted guidance notes are included following each score level to support consistent interpretation of the criteria and to resolve potential ambiguities during scoring. This layered design; combining rubric-based descriptors, explicit performance features, illustrative examples, and practical annotation notes; aims to enhance clarity, reduce subjectivity, and improve inter-rater consistency by operationalizing abstract scoring criteria into interpretable and actionable annotation decisions.

**Annotated Arabic Examples** The guidebook includes multiple expert-annotated Arabic essays at each score level, covering both text types. Each example provides the complete Arabic text, all seven trait scores, and Arabic-language justifications with highlighted key features. Figure 1 in the Appendix shows a real annotated essay from the dataset for an expository prompt (P7: “Staying Up Late”), illustrating how the scoring descriptors are applied in practice.

**Edge Case Protocols** The guidebook provides explicit protocols for: **copied or AI-generated text** (score as written); **code-switching** between MSA and colloquial Arabic or English (score MSA portions; heavy non-Arabic content may affect REL and VOC); **very short essays** (<50 words; score on what is present); and **error-dense text** (score on what can be understood). Annotators are reminded that traits are independent: correct Arabic spelling does not compensate for weak content.

A recurring challenge is the use of Arabic enumeration markers (أولاً، ثانياً، وأخيراً) alone. Through calibration, guidelines were clarified to specify that ORG requires both structural elements *and* meaningful thematic connections, as shown in the Score 5 justification in Figure 1.

### 4.3 Addressing Subjectivity

We employed four Arabic-specific strategies to minimize annotator bias: (1) evaluate each trait independently to avoid halo effects from ornate Arabic style; (2) avoid over relying on essay length or classical Arabic expressions in isolation; (3) apply opposite standards for DEV depending on text type, rewarding objectivity in expository and position-taking in persuasive; and (4) provide diverse Arabic exemplars at each score level to illustrate multiple paths to the same score.

## 5 Annotator Training Protocol

Our five-stage protocol was implemented over four weeks before production annotation began.

**Stage 1 – Conceptual Training (Week 1):** Joint review of the Arabic guidebook, discussion of text-type distinctions with Arabic examples, examination of trait descriptors, and group exercises identifying Arabic features in sample essays.

**Stage 2 – Expert Model Exposure (Week 2):** Review of 8 expert-annotated Arabic essays (4 expository, 4 persuasive). Supervisors explained score assignments; discussion focused on adjacent score levels and borderline cases, e.g., distinguishing surface connectors (أدوات الربط) from genuine structural cohesion.

**Stage 3 – Independent Practice (Week 3):** Each annotator independently scored 12 Arabic practice essays (6 per text type) without access to expert scores, generating diagnostic data on individual understanding of Arabic-specific criteria.

Descriptor	Performance Characteristics	Illustrative Example	Score
Intro and conclusion absent. No organization or logical sequence. Ideas random or disconnected.	No clear introduction or conclusion. The text consists of a single paragraph or several unstructured sentences. Ideas are presented without sequence or connectors.	The student begins directly with: "People buy things online" (الناس يشترون من النت), then lists benefits in a scattered manner, with no ending.	1
Either intro or conclusion absent. Paragraphs lack logical progression. Weak organization attempt.	Only one of the introduction or conclusion is present; the other is completely absent or non-functional. The body consists of weakly connected or non-sequential paragraphs.	The student writes a general introduction such as: "Online shopping is beneficial" (الشراء الإلكتروني مفيد), then presents one or two ideas without clear connection or progression, ending with an incomplete phrase such as: "and so on" (وهكذا).	2
Both intro and conclusion present, but 1–2 body paragraphs lack coherence or connection to main idea.	Both introduction and conclusion are present, but not highly effective. The body includes one or two paragraphs with some coherence, though sequencing is weak or sometimes absent.	The introduction presents the topic in general terms; the paragraphs discuss two reasons for the phenomenon, but the connection between them is unclear; the conclusion repeats the initial idea without drawing a conclusion.	3
Appropriate intro and conclusion. 2–3 sequential, coherent body paragraphs. Minor transition issues only.	The introduction effectively introduces the topic and presents the main idea. The body is divided into clear paragraphs with logical progression and gradual development of ideas. The conclusion is appropriate and summarizes the points without repetition.	The student begins with an introduction such as: "With technological advancement, e-commerce has become a necessity" (مع تطور التكنولوجيا، أصبحت التجارة الإلكترونية ضرورة) then presents multiple reasons, each in a separate paragraph, and concludes with a sentence emphasizing the importance of the phenomenon.	4
Effective intro, strong conclusion, 2–3 body paragraphs with clear topic sentences and smooth transitions.	The introduction is direct and engaging (e.g., starts with a question, fact, or quotation). The body presents fully developed and interconnected ideas using cohesive devices. The conclusion provides a general insight or recommendation that highlights the discussion.	The student begins with: "Can we imagine a world without online stores?" (هل يمكن أن تتصور عالماً بلا متاجر إلكترونية؟) develops paragraphs explaining the reasons for its spread in an organized manner, and concludes with: "Therefore, e-commerce is no longer an option, but a necessity" (لذلك، فإن التجارة الإلكترونية) بل ضرورة). لم تعد خياراً، بل ضرورة	5

Table 2: Organization (ORG) scoring descriptors (translated from Arabic).

#### Stage 4 – Calibration and Discussion (Week 4):

Calculation of individual agreement with expert scores; group discussion of high-disagreement Arabic cases; re-scoring of 4 essays together. A concrete calibration finding was the early confusion about enumeration markers: annotators initially awarded  $ORG = 5$  for essays using "first, second, finally" (أولاً، ثانياً، أخيراً) without thematic paragraph development; calibration established that such essays should receive  $ORG = 3-4$  depending on paragraph coherence.

**Stage 5 – Monitored Production (ongoing):** Annotation production with dual annotation, weekly calibration sessions on challenging cases, and random control essay insertion (see Section 6).

#### 5.1 Annotators

The annotation team consisted of six annotators and three supervisors. All members of the team were Arabic language teachers or lecturers with formal training in Arabic language education. Five members of the team held advanced degrees (MSc or PhD) in Arabic language or linguistics.

Annotators were responsible for the primary essay scoring tasks, while supervisors oversaw annotator training, quality assurance procedures, and dispute resolution during the annotation process. Supervisors also conducted periodic calibration sessions and monitored annotation quality through control essays and adjudication reviews.

## 5.2 Annotation Guidelines

All essays in the dataset were evaluated using the Core Academic Skills Test rubric developed by the Qatar University Testing Center (QUTC)<sup>1</sup>. The rubric evaluates seven writing traits: Relevance (REL), Organization (ORG), Vocabulary (VOC), Style (STY), Development (DEV), Mechanics (MEC), and Grammar (GRA). In addition, a Holistic score (HOL) was computed as the sum of the individual trait scores.

Six traits (ORG, VOC, STY, DEV, MEC, GRA) were rated on a 6-point scale (0 = lowest, 5 = highest), while Relevance (REL) was rated on a 3-point scale (0 = not relevant, 1 = partially relevant, 2 = fully relevant). If an essay received a REL score of 0, all remaining trait scores were automatically set to 0, since responses that do not address the prompt are not subject to further evaluation.

To ensure consistent interpretation of the rubric, two supervisors developed a comprehensive annotation guidebook containing detailed scoring terminology, annotated examples, and practice exercises for each prompt type.<sup>2</sup> Annotators were required to review the guidebook and complete structured training sessions before beginning annotation production.

Following training, moderation sessions were conducted in which annotators jointly reviewed sample essays, discussed scoring discrepancies, and harmonized interpretations of the rubric. These sessions ensured consistent application of the rubric across the annotation team before large-scale annotation began.

## 6 Quality Assurance Mechanisms

To ensure annotation reliability, we implemented several quality assurance procedures supervised by the senior annotation team.

**Dual Annotation.** Every essay was independently scored by two annotators (R1 and R2) under blind conditions using the Assessment Gourmet Platform<sup>3</sup>, which anonymized student identity and prevented annotators from seeing each other’s scores.

**Discrepancy Resolution.** If the difference between the two holistic scores (HOL) was less than

<sup>1</sup>[https://www.qu.edu.qa/sites/en\\_US/testing-center/TestDevelopment/cast](https://www.qu.edu.qa/sites/en_US/testing-center/TestDevelopment/cast)

<sup>2</sup>The full annotation guidelines are available at: [https://gitlab.com/bigirqu/laila/-/raw/main/rubrics/annotation\\_guidebook.pdf](https://gitlab.com/bigirqu/laila/-/raw/main/rubrics/annotation_guidebook.pdf)

<sup>3</sup><https://g-assess.com/>

P#	Type	Essays	Avg Len	R3%
P1 Sports	EXP	1,122	162	10.4
P2 Social Media	PER	1,168	175	15.5
P3 Technology	PER	521	159	9.4
P4 Communication	PER	500	152	15.0
P5 Heritage	EXP	1,181	157	23.3
P6 Homework	PER	1,162	160	20.3
P7 Staying Up	EXP	1,143	202	10.6
P8 Video Games	PER	1,062	186	15.7
<b>Total</b>		<b>7,859</b>	<b>171</b>	<b>15.7</b>

Table 3: Dataset statistics. EXP/PER = expository/persuasive. R3% =percentage of essays requiring supervisor adjudication.

6 points (approximately 19% of the maximum possible score of 32), the scores were averaged. Larger discrepancies ( $\geq 6$  points) were escalated to a supervising annotator (R3), who performed adjudication and provided written feedback to the original annotators. This process served both as conflict resolution and as an ongoing learning mechanism.

Overall, 15.7% of essays required supervisor adjudication (range: 10.4%–23.3% depending on the prompt), as reported in Table 3.

**Control Essay Monitoring.** Approximately 5% of the essays assigned to each annotator were pre-scored expert control essays inserted without the annotator’s knowledge. Inter-annotator agreement was monitored using Quadratic Weighted Kappa (QWK). If an annotator’s QWK dropped below 0.60 over any 20-essay window, annotation was paused and additional calibration training was conducted with supervisors.

## 7 Results

We report results across two dimensions: the operational characteristics of the annotation process itself, and the inter-annotator agreement achieved under the finalized methodology.

### 7.1 Application Context

The methodology was applied to build an Arabic Automatic Essay Scoring dataset (Bashendy et al., 2025b): 7,859 essays by 4,372 high school students across 8 prompts (3 expository, 5 persuasive), annotated over one academic year by 6 annotators under 3 senior supervisors. Table 3 summarizes prompt-level statistics.

**Adjudication patterns.** Overall, 15.7% of essays required supervisor adjudication (R3), ranging from 10.4% (P1,Sports) to 23.3% (P5, Her-

Trait	P1	P2	P3	P4	P5	P6	P7	P8
REL	.79	.60	.67	.59	.58	.68	.75	.77
ORG	.78	.72	.83	.77	.74	.78	.77	.78
VOC	.74	.71	.76	.69	.71	.75	.79	.80
STY	.74	.71	.76	.72	.60	.71	.68	.72
DEV	.79	.64	.58	.66	.70	.72	.71	.55
MEC	.69	.65	.77	.72	.58	.65	.69	.70
GRA	.71	.65	.77	.70	.70	.73	.67	.61
Avg	<b>.75</b>	<b>.67</b>	<b>.73</b>	<b>.69</b>	<b>.66</b>	<b>.72</b>	<b>.72</b>	<b>.70</b>

Table 4: Inter-annotator agreement (QWK) by prompt and trait. The majority of values fall in the substantial range (0.61–0.80); moderate values (<0.61) occur in nine prompt–trait cells, concentrated in REL and DEV across semantically open prompts (P2, P4, P5).

itage). This variation is not random: adjudication rates correlate with prompt framing breadth. Narrowly framed prompts with concrete phenomena (P1 Sports, P7 Staying Up Late) produced the lowest adjudication rates (10.4% and 10.6%), as the essay topic constrained valid interpretations of REL and DEV. Broadly framed prompts admitting diverse valid positions (P5 Heritage: “*balance between heritage and modernity*”) drove higher disagreement, confirming that prompt design is a significant but often overlooked source of annotation variance. We recommend pilot-testing prompts on a small essay sample and targeting adjudication rates below 15% as an indicator of adequate prompt specificity before large-scale annotation.

## 7.2 Inter-Annotator Agreement

We measured agreement between R1 and R2 (before adjudication) using Quadratic Weighted Kappa (QWK) following (Landis and Koch, 1977): <0.40 (poor), 0.40–0.60 (moderate), 0.61–0.80 (substantial), 0.81–1.00 (almost perfect). QWK is particularly appropriate for ordinal scales such as ours, as it penalizes disagreements proportionally to their distance on the scale, making it more sensitive than simple percent agreement for detecting systematic annotator bias.

Table 4 shows that the methodology achieved substantial agreement ( $\geq 0.61$ ) in the large majority of prompt–trait combinations, with an overall average QWK of 0.71. Nine cells fall into the moderate range (<0.61): these are not distributed randomly but cluster in two traits (REL and DEV) and three prompts (P2, P4, P5), all of which involve persuasive writing on open-ended social topics. This pattern is theoretically coherent: both REL and DEV require annotators to exercise semantic judg-

ment about argument quality and topical relevance, precisely the dimensions most sensitive to prompt framing.

Across traits, ORG achieved the highest average agreement (0.77), with no prompt falling below 0.72. This is consistent with our design decision to anchor ORG descriptors in structurally observable Arabic features; presence of introduction and conclusion, paragraph count, and explicit topic sentences; which leave less room for subjective interpretation. At the other end, DEV showed the highest variance across prompts (range: 0.55–0.79), reflecting the inherent difficulty of operationalizing Arabic argumentation quality, where annotators must assess not only what claims are made but whether evidence is culturally appropriate and rhetorically effective by MSA norms. The low DEV agreement on P8 (Video Games, 0.55) and P3 (Technology, 0.58) warrants attention: both are persuasive prompts on technology topics where student essays frequently blurred the boundary between anecdotal opinion and substantiated argument, a distinction our guidelines address but which remains challenging to apply consistently.

REL shows the widest prompt-level range of any trait (0.58–0.79), a finding that directly informs prompt design: prompts receiving lower REL agreement (P4: 0.59, P5: 0.58) are precisely those with broader, more abstractly framed topics, where the threshold for “partial relevance” (REL = 1) versus “full relevance” (REL = 2) is harder to determine. This suggests that REL agreement is as much a function of prompt specificity as of guideline quality, and future work should consider tightening prompt framing as a complementary strategy to guideline refinement.

## 7.3 Agreement Patterns

**Trait-level variation.** ORG achieved the highest average agreement (0.77), with no prompt falling below 0.72. This consistency reflects our design decision to anchor ORG descriptors in structurally observable Arabic features; presence of introduction and conclusion, paragraph count, and explicit topic sentences; which can be verified directly against the essay text and generalize well across both text types. DEV showed the lowest average agreement (0.66) and the highest variance across prompts (range: 0.55 - 0.79), reflecting the inherent subjectivity of judging Arabic argumentation quality - consistent with the prior educational assess-

ment literature (Williamson et al., 2012). The two lowest individual DEV values occur on persuasive prompts (P8: 0.55; P3: 0.58), where student essays frequently blurred the boundary between anecdotal opinion and substantiated argument. The correlation between DEV agreement and adjudication rate further supports this: prompts with higher R3% (e.g., P5 at 23.3%) consistently show lower DEV agreement in Table 3, suggesting that semantic trait subjectivity and prompt ambiguity are compounding rather than independent sources of annotation variance.

**Prompt-level variation.** Agreement ranged from 0.66 (P5 Heritage) to 0.75 (P1 Sports). The Heritage prompt’s broad framing (“balance between heritage and modernity”) admits diverse valid interpretations, making REL and DEV judgments harder and driving a 23.3% adjudication rate versus the project average of 15.7%.

**Semantic vs. surface traits.** Traits requiring deeper Arabic semantic understanding (DEV, REL) show lower agreement than surface-level traits (MEC, GRA), consistent with the challenge of codifying Arabic argumentative quality.

**Learning curve.** During the early stages of annotation, substantial variability was observed across annotators, reflecting differences in interpreting the scoring criteria. Following calibration, this variability decreased, and agreement improved further with continued practice. However, variability re-emerged with the introduction of new prompts, indicating the need for renewed alignment when encountering unfamiliar responses. This pattern is consistent with previous findings on the transition from guided calibration to independent scoring of unseen responses (Williamson et al., 2012). Moderation sessions between the supervisor and individual annotator helped restore consistency. Overall, these observations suggest that annotator agreement is affected by both task familiarity and the progressive clarification of the guidelines, with the final, post-calibration version supporting substantially higher and more stable agreement.

## 8 Discussion

**Key success factors.** Five elements proved most critical: (1) an Arabic-language guidebook with concrete student writing examples, especially for STY where abstract cohesion definitions required grounding; (2) the terminology unification section,

which resolved early disagreements from inconsistent interpretation of Arabic linguistic terms; (3) text-type awareness training that calibrated opposite scoring mindsets for expository versus persuasive writing; (4) multi-stage training with calibration before production began; and (5) continuous quality assurance throughout the year-long annotation.

**Challenges.** The greatest challenge was semantic trait subjectivity: DEV and REL consistently yielded lower IAA than structural traits. We addressed this through extensive Arabic exemplars and evidence-based reasoning during training, but some residual subjectivity is unavoidable. Future work might decompose DEV into sub-traits (claim quality, evidence quality, counterargument handling). A second challenge was the steep learning curve when annotators first used the guidebook, especially for semantic traits like DEV and REL. Early calibration showed that some disagreements came from residual ambiguities in certain descriptors and edge-case protocols, not just trait subjectivity. We therefore refined the wording of these sections (for example, distinguishing superficial enumeration from genuine paragraph development in ORG and clarifying the threshold between partial and full relevance in REL) and expanded the Arabic borderline examples. These revisions increased and stabilized inter-annotator agreement across prompts, suggesting that the final guidelines were sufficiently clear despite the initial learning curve. Prompt design was also critical: the Heritage prompt’s broad framing drove significantly higher adjudication (23.3%), informing our recommendation to pilot-test prompts before large-scale annotation.

**Generalizability.** While developed for MSA in a specific educational context, the core components - operational definitions in the target language, native-language exemplars, multi-stage training, dual annotation, continuous calibration - apply universally. Researchers adapting this approach to other Arabic varieties, proficiency levels, or languages should maintain these core components; shortcuts in any area will likely compromise annotation quality.

## 9 Conclusion and Future Work

We presented a systematic methodology for developing annotation guidelines for trait-based Arabic

AES, achieving substantial inter-annotator agreement (QWK: 0.66–0.75) across 7,859 Arabic essays. Our results demonstrate that reliable large-scale Arabic essay annotation is achievable when scoring criteria are operationalized in observable, Arabic-specific terms, annotators undergo structured calibrated training, and quality assurance is maintained continuously, evidenced by agreement rising before production began. Beyond the dataset, we offer a replicable template whose core components (target-language definitions, native exemplars, multi-stage calibration, and dual annotation) generalize to other morphologically complex, under-resourced languages. Several directions follow from the limitations identified above. First, decomposing DEV into sub-traits (claim quality, evidence quality, counterargument handling) may improve both reliability and model signal, given its high agreement variance (0.55–0.79). Second, our adjudication data suggest that prompt specificity is quantifiable; developing explicit prompt design criteria to target adjudication rates below 15% before deployment is a practical next step. Third, extending the framework to other Arabic varieties, narrative and descriptive genres, and university-level writing would broaden its applicability. Finally, applying AES models for automated pre-annotation could reduce annotator load while preserving reliability, particularly for surface traits where agreement is already high. Remaining challenges, particularly in semantic traits (DEV, REL), point to prompt design as a variable that guideline refinement alone cannot resolve. All materials are publicly released to lower the barrier for future annotation efforts in Arabic and comparable languages.

## Limitations

**Arabic variety:** Guidelines target MSA in a specific educational context; adaptation may be required for other Arabic varieties. **Genre coverage:** Only expository and persuasive writing are addressed; narrative and descriptive Arabic writing require separate criteria. **Proficiency range:** Framework targets grades 10–12; university or younger students require recalibration. **Resource intensity:** Four weeks of training plus ongoing calibration may be prohibitive for low-budget contexts. **Essay length:** Essays averaged 171 words; effectiveness for longer compositions remains to be demonstrated.

## Ethical Considerations

All essays were collected from high school students; participation was voluntary with consent obtained from students and guardians following institutional guidelines.<sup>4</sup> No personally identifying information appears in the released dataset. Annotators were compensated fairly and workloads were monitored to prevent fatigue. The dataset is released for research use only.

## Acknowledgments

We heartily thank our dedicated annotators for their contributions and express our gratitude to Qatar University Testing Center, the Ministry of Education and Higher Education in Qatar (the Arabic Section of the Department Of Educational Supervision in particular). This work was made possible by NPRP grant NPRP14S-0402-210127 from the Qatar Research Development and Innovation (QRDI) Council. The statements made herein are solely the responsibility of the authors.

## References

- Ron Artstein and Massimo Poesio. 2008. Inter-coder agreement for computational linguistics. *Computational Linguistics*, 34(4):555–596.
- May Bashendy, Salam Albatarni, Sohaila Eltanbouly, Walid Massoud, Houda Bouamor, and Tamer Elsayed. 2025a. TAQEEM 2025: Overview of the first shared task for Arabic quality evaluation of essays in multidimensions. In *Proceedings of the Third Arabic Natural Language Processing Conference (ArabicNLP 2025)*. To appear.
- May Bashendy, Salam Albatarni, Sohaila Eltanbouly, Eman Zahran, Hamdo Elhuseyin, Tamer Elsayed, Walid Massoud, and Houda Bouamor. 2024. QAES: First publicly-available trait-specific annotations for automated scoring of Arabic essays. In *Proceedings of the Second Arabic Natural Language Processing Conference*, pages 337–351, Bangkok, Thailand. Association for Computational Linguistics.
- May Bashendy, Walid Massoud, Sohaila Eltanbouly, Salam Albatarni, Marwan Sayed, Abrar Abir, Houda Bouamor, and Tamer Elsayed. 2025b. LAILA: A large trait-based dataset for Arabic automated essay scoring. In *Proceedings of the 2025 Conference of the North American Chapter of the Association for Computational Linguistics*. To appear.
- Adriane Boyd, Jirka Hana, Lionel Nicolas, Detmar Meurers, Katrin Wisniewski, Andrea Abel, Karin Schöne, Barbora Stindlová, and Chiara Vettori. 2014.

<sup>4</sup>IRB Number: QU-IRB 159/2024-EA)

The MERLIN corpus: Learner language and the CEFR. In *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC'14)*, pages 1281–1288, Reykjavik, Iceland.

Scott Crossley, Yu Tian, Perpetual Baffour, Alex Franklin, Youngmeen Kim, Wesley Morris, Meg Benner, Aigner Picou, and Ulrich Boser. 2023a. The English language learner insight, proficiency and skills evaluation (ELLIPSE) corpus. *International Journal of Learner Corpus Research*, 9(2):248–269.

Scott Andrew Crossley, Perpetual Baffour, Yu Tian, Alex Franklin, Meg Benner, and Ulrich Boser. 2023b. A large-scale corpus for assessing written argumentation: PERSUADE 2.0. Available at SSRN 4795747.

Karen Fort, Gilles Adda, and K. Bretonnel Cohen. 2011. Amazon mechanical turk: Gold mine or coal mine? volume 37, pages 413–420.

Nizar Habash and David Palfreyman. 2022. ZAEBUC: An annotated Arabic-English bilingual writer corpus. In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 79–88, Marseille, France. European Language Resources Association.

Yaqiong He, Feng Jiang, Xiaomin Chu, and Peifeng Li. 2022. Automated Chinese essay scoring from multiple traits. In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 3007–3016, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.

J. Richard Landis and Gary G. Koch. 1977. The measurement of observer agreement for categorical data. *Biometrics*, 33(1):159–174.

Sandeep Mathias and Pushpak Bhattacharyya. 2018. Thank you for attending!: Attending to latent representation and modeling textual coherence in Automated Essay Scoring. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*.

Rodrigo Wilkens, Alice Pintard, David Alfter, Vincent Folny, and Thomas François. 2023. TCFLE-8: A corpus of learner written productions for French as a foreign language and its application to automated essay scoring. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 3447–3465, Singapore. Association for Computational Linguistics.

David M. Williamson, Xiaoming Xi, and F. Jay Breyer. 2012. A framework for evaluation and use of automated scoring. *Educational Measurement: Issues and Practice*, 31(1):2–13.

## A Full Development Trait Descriptors

### Expository Essays (النص التفسيري)

**Score 1:** Content largely unrelated to the topic. Ideas random, incoherent, no logical sequence.

Main idea absent. No analysis. **Score 2:** Content somewhat related. Main idea disappears. Limited topic coverage. May contain personal opinion inappropriate for expository Arabic writing. **Score 3:** Completely on topic. Ideas mostly sequential but main idea fades. Some evidence but disorganized. Expository neutrality maintained but analytical depth insufficient. **Score 4:** Completely on topic. Ideas clear, organized, coherent. Main idea consistently connected to sub-ideas. Specific explanations and coherent supporting evidence, though not comprehensive. **Score 5:** Completely on topic throughout. Main idea strongly maintained. Comprehensive explanations. Multiple evidence forms. Full expository neutrality (حياد وموضوعية) maintained.

### Persuasive Essays (النص الإقناعي)

**Score 1:** No clear position taken. Ideas random. No arguments or evidence. Persuasive techniques (الأساليب الإقناعية) entirely absent. **Score 2:** Position unclear or inconsistent. Limited argumentation, no evidence. Does not engage with alternative viewpoints. **Score 3:** Position stated but not consistently maintained. Arguments underdeveloped. Some evidence, not well-integrated. **Score 4:** Clear position adopted and maintained. Solid arguments with evidence. Some persuasive techniques. May not fully address counterarguments. **Score 5:** Strong, clear position throughout. Well-developed arguments with robust evidence (facts, examples, quotations, statistics). Effective persuasive techniques. Explicitly acknowledges and refutes counterarguments (عرض الآراء المختلفة والرد عليها).

## B Example

FIGURE A.1 Annotated Expository Essay — Prompt P7: "Staying Up Late" (السهر) Type: Expository · Student: HS Gr. 11

TRAIT SCORES REL 2/5 ORG 5/5 VOC 4/5 STY 4/5 DEV 4/5 MEC 5/5 GRA 4/5 HOL 28/32

RELEVANCE (REL) ● Fully relevant — essay addresses all aspects of the prompt

STUDENT ESSAY TEXT (WITH ANNOTATION HIGHLIGHTS)

PROMPT  
اكتب مقالة تفسيرية تتناول فيها ظاهرة السهر لدى الطلاب وأسبابها وتأثيراتها.

يُعدّ السهر من أكثر العادات انتشاراً في المجتمعات الحديثة، ولا سيّما في أوساط طلاب المدارس والجامعات. <sup>1</sup> وتُشير الدراسات إلى أن نسبة كبيرة من الطلاب يسهرون حتى ساعات متأخرة من الليل. <sup>2</sup> مما ينعكس سلباً على صحتهم وتحصيلهم الدراسي. وفي ضوء ذلك، تسعى هذه المقالة إلى تناول هذه الظاهرة بالتحليل والتفسير الموضوعي. <sup>3</sup>

ولعلّ من أبرز أسباب السهر <sup>1</sup> استخدام وسائل التواصل الاجتماعي والأجهزة الإلكترونية؛ إذ بات الهاتف المحمول ريفاً دائماً للشباب في ساعات الليل. <sup>4</sup> فضلاً عن ذلك، تُسهم الضغوط الدراسية والامتحانات في إطالة أوقات اليقظة، <sup>5</sup> حيث يضطرّ كثير من الطلاب إلى المراجعة في أوقات متأخرة لاستيعاب المادة العلمية. <sup>6</sup>

إنّما تأثيرات السهر المتكرر فهي متنشعبة وخطيرة. <sup>1</sup> على الصعيد الصحي، يُؤدي قلّة النوم إلى ضعف التركيز وتراجع المناعة وارتفاع مستويات التوتر. <sup>2</sup> وأما على الصعيد الأكاديمي، فقد تبين أن الطلاب الذين يحطون بقدرٍ كافٍ من النوم يُظهرون أداءً أكاديمياً أفضل بكثير ممّن يعانون من الحرمان منه. <sup>3</sup> وتُدرِّم هذه النتيجة دراسة المعهد الوطني للصحة التي أثبتت أن النوم يُحسّن الذاكرة طويلة الأمد. <sup>4</sup>

خلاصة القول، إنّ ظاهرة السهر مسألة تستدعي الاهتمام والمعالجة الجادة. <sup>5</sup> وتقع المسؤولية على عاتق الأسرة والمؤسسات التعليمية في توعية الطلاب بأهمية الالتزام بمواعيد النوم المنتظمة، <sup>6</sup> واتخاذ خطوات عملية للحدّ من الآثار السلبية لهذه الظاهرة على الأجيال القادمة. <sup>7</sup>

ANNOTATION NOTES

● ORG — ORGANIZATION 5/5

1 مقدمة مباشرة وقالة تطرح الظاهرة وتحدد محور المقالة. كل فقرة تحتوي على جملة موضوعية واضحة (السهر – أسبابه – آثاره). الخاتمة تلخص وتقدم توصية دون تكرار. استخدمت أدوات ربط بين الفقرات: أما، فضلاً، خلاصة القول.

● VOC — VOCABULARY 4/5

2 مفردات متنوعة ودقيقة: تشعبية، الحرمان، الأمد البعيد، غياب الحشو. يُخصم نقطة واحدة لغياب الصيغيات الاصطلاحية الأرقى وبعض التكرار المعجمي في الفقرة الثانية.

● STY — STYLE & COHESION 4/5

1 ربط خطي فقال: إجاله هذه (الظاهرة)، حذف، وأدوات وصل متنوعة. ربط معجمي من خلال التضاد (اليقظة/ النوم) والحفل الدلالي. يُخصم لغياب أنماط التنظيم الأسلوبية الأعمق.

● DEV — DEVELOPMENT 4/5

2 يعالج الأسباب والآثار بموضوعية تفسيرية، يستشهد بمصدر بحثي (المعهد الوطني للصحة). يُخصم لعدم الاستشهاد بأكثر من مصدر، وغياب الأرقام الكمية الداعمة.

● MEC — MECHANICS 5/5

1 تشكيل صحيح، همزات سليمة، ناء مربوطة لا ليس فيها، ترقيم مناسب. لا أخطاء إملائية مرصودة.

● GRA — GRAMMAR 4/5

1 تنوع تركيب جمل شرطية ضمنية، تراكيب اسمية وفعليّة. يُخصم لغياب الحفل الوصفية والمركبات الاعتراضية المتنوعة في الكتابة الأكاديمية المتقدمة.

HIGHLIGHT KEY ORG VOC STY DEV MEC GRA

Figure 1: Expert-annotated expository student essay (P7: “Staying Up Late”, Grade 11). Colored highlights link textual evidence to the six scored traits; circled superscript numbers correspond to annotation notes in the right panel. REL = 2/2 (fully relevant); HOL = 28/32. Justifications are rendered in Modern Standard Arabic, matching the annotation protocol. ORG achieves the maximum score (5/5) owing to structurally observable features: a direct, engaging introduction; three body paragraphs each with an explicit topic sentence; and a conclusion offering a recommendation rather than mere repetition—consistent with the Score 5 descriptor in Table 2. The one-point deductions in VOC, STY, DEV, and GRA reflect specific, evidence-cited shortfalls rather than global impressions, operationalizing the evidence-based scoring principle.

TRAIT SCORES **REL 2 /2** **ORG 5 /5** **VOC 4 /5** **STY 4 /5** **DEV 4 /5** **MEC 5 /5** **GRA 4 /5** **HOL 28 /32**

**RELEVANCE (REL)** ●●○ Fully relevant — essay addresses all aspects of the prompt

STUDENT ESSAY TEXT (WITH ANNOTATION HIGHLIGHTS)

PROMPT  
Write an expository essay in which you discuss the phenomenon of students staying up late, its causes, and its effects.

Staying up late is considered one of the most widespread habits in modern societies, especially among school and university students. 1 Studies indicate that a large percentage of students stay up until late hours of the night, which is reflected negatively on their health and academic achievement. 2 In light of that, this essay seeks to address this phenomenon through objective analysis and explanation. 4

Perhaps one of the most prominent causes of staying up late is the use of social media and electronic devices; the mobile phone has become a permanent companion for young people during the night hours. In addition to that, academic pressures and exams contribute to extending the times of wakefulness, 5 as many students are forced to review at late times in order to absorb the scientific material. 3

As for the effects of repeated staying up late, they are branched and serious. 1 On the health level, lack of sleep leads to weak concentration, decline in immunity, and high levels of stress. As for the academic level, it has become clear that students who get a sufficient amount of sleep show much better academic performance than those who suffer from deprivation of it. This result is supported by a study by the National Institutes of Health, which proved that sleep improves long-term memory. 2

In conclusion, the phenomenon of staying up late is an issue that calls for serious attention and treatment. 1 Responsibility falls on the family and educational institutions to make students aware of the importance of committing to regular sleep times and taking practical steps to limit the negative effects of this phenomenon on future generations. 3

ANNOTATION NOTES

**ORG — ORGANIZATION** 5 /5  
1 A direct and effective introduction presents the phenomenon and identifies the focus of the essay. Each paragraph contains a clear topic sentence: staying up late → its causes → its effects. The conclusion summarizes the discussion and provides a recommendation without repetition. Connective devices are used between paragraphs, such as "as for," "in addition," and "in conclusion."

**VOC — VOCABULARY** 4 /5  
5 The essay uses varied and precise vocabulary, such as "branched," "deprivation," and "long-term." There is no filler language. One point is deducted because the essay lacks more advanced idiomatic phrasing and includes some lexical repetition in the second paragraph.

**STY — STYLE & COHESION** 4 /5  
4 The essay demonstrates effective linear cohesion, including reference, as in "this phenomenon," ellipsis, and varied connective devices. It also shows lexical cohesion through contrast, such as "wakefulness / sleep," and through semantic field relationships. One point is deducted because deeper stylistic organizational patterns are absent.

**DEV — DEVELOPMENT** 4 /5  
2 The essay addresses causes and effects with expository objectivity. It cites a research source, the National Institutes of Health. One point is deducted because the essay does not cite more than one source and lacks supporting quantitative evidence.

**MEC — MECHANICS** 5 /5  
5 The text shows correct orthography and appropriate punctuation throughout. No spelling errors are observed.

**GRA — GRAMMAR** 4 /5  
3 The essay demonstrates good syntactic variety, including implicit conditional structures and both nominal and verbal sentence structures. One point is deducted because the essay lacks descriptive clauses and parenthetical constructions expected in more advanced academic writing.

HIGHLIGHT KEY — **ORG** **VOC** **STY** **DEV** **MEC** **GRA**

Figure 2: English translation of Figure 1: annotated expository essay with scoring rationale for each trait.

# Revisiting Faithfulness Annotations for Long-form Summaries

Yang Zhong<sup>1</sup>, Yang Janet Liu<sup>\*2</sup>, Diane Litman<sup>\*1,3</sup>

<sup>1</sup>Department of Computer Science, <sup>2</sup>Department of Linguistics

<sup>3</sup>Learning Research & Development Center

University of Pittsburgh, Pittsburgh, PA, USA

{yaz118, jal787, dlitman}@pitt.edu

## Abstract

Benchmarks for long-form summaries (four or more sentences) generated by language models increasingly serve as gold-standard references for developing, evaluating, and comparing faithfulness-checking systems. As their influence grows, understanding the challenges of annotating faithfulness errors within long, discourse-rich summaries becomes critical. We revisit three benchmarks spanning diverse text types and contrasting annotation designs. Using a discourse-aware evaluation framework together with human auditing, we identify cases where benchmark labels may be unreliable. Manual verification shows that 3.4%-5.4% of sentence-level labels warrant revision due to discourse-level inconsistencies that standard annotation procedures overlook. We introduce a taxonomy of five recurring annotation error types, propose revised labels, and show that correcting these cases leads to meaningful shifts in system rankings. We conclude with recommendations for future annotation practices.

## 1 Introduction

Faithfulness benchmarks for AI-generated summaries have become de facto gold-standard references for developing and comparing automatic fact-checking systems (Kryscinski et al., 2019; Fabri et al., 2021; Tang et al., 2023), making their annotation reliability a foundational concern. Evaluating long-form summaries presents distinct challenges. Unlike news-domain benchmarks, where summaries typically contain one to three sentences, long-form summaries often exceed 100 words and synthesize information from multiple parts of the source document. This increased complexity requires annotators to assess not only the accuracy of individual facts but also the coherence and validity of relationships among them. Prior work has shown that such demands impose substantial cognitive load and can reduce inter-annotator agreement

\*Equal senior contribution.

**Source Document**  
[Abstract] ... We report the case of a 53-year-old woman presented with ... extensive vascular nevus, which match the typical manifestations of phakomatosis pigmentovascularis of cesioflammea type, according to Happle's classification. **The rare occurrence of this genodermatosis and the clinical exuberance of the skin lesions motivated this case report. ...**

**Summary**  
... The patient's case is **notable because** it matches the typical manifestations of phakomatosis pigmentovascularis ... as described by a doctor named Happle.

**Initial Annotation**  
Label: Faithful  
Marked Evidence: We report the case ... which match the manifestation ... according to Happle's classification.  
**Key issue:** The annotator did not verify the underlying reason supporting the claim that the case is notable.

How do we check for potential annotator oversights?

**This Work**  
We use a discourse-aware evaluation framework to check the summaries, then manually verify the predicted labels and rationales to detect and analyze annotator errors.

**Discourse-aware Framework**  
Prediction: Unfaithful

**Discourse-based Decomposition**  
C9: The case matches the typical manifestations of ... as described by Dr. Happle.  
C10: The 53-year-old woman's case is notable.  
Relations: (C9, C10, Causal)

**Model Rationale**  
Source states the report was motivated by 'rare occurrence' and 'clinical exuberance,' not 'because it matches typical manifestations.'. **The causal link (match → notable) in the summary is not established in the source (C9→C10 causal is fabricated)**

Figure 1: **Human annotators overlooked discourse-level errors.** The source document attributes the case report's motivation to "the rare occurrence" and the "clinical exuberance". However, the annotator (a domain expert) relied on a preceding evidence sentence and failed to evaluate the *causal relationship* underlying the case's significance.

(Krishna et al., 2023). Moreover, faithfulness errors in long-form summaries increasingly arise from subtle distortions in meaning or framing (Zhong and Litman, 2025b), rather than from surface-level inconsistencies (Goyal et al., 2022).

A closer examination of long-form summaries reveals that many problematic cases stem not from incorrect individual facts, but from whether the *relationships* between them are sufficiently established. For example, Figure 1 shows that a human annotator labeled a summary sentence as faithful because the individual facts appear in the source. However, the annotation overlooks whether the causal relationship—namely, that "the matching of the description of a patient's case" makes it "notable"—is actually supported by the document. Such errors

are difficult to detect when annotation protocols prioritize verifying the presence of discrete facts over assessing how those facts are connected. Similar issues arise when summaries alter temporal ordering, omit hedging that overstate certainty, or recast factual content with interpretive language.

Motivated by these challenges, we revisit three recent benchmarks for evaluating faithfulness in long-form summaries, which span diverse domains and annotation designs. Using a discourse-aware evaluation framework together with complementary LLM-based analysis methods, we conduct a systematic audit of sentences where automatic evaluators disagree with the gold annotations. For these cases, we perform manual verification to assess whether the original labels remain justified (§4).

Our reassessment reveals that 3.4–5.4% of sentence-level annotations require revision upon closer examination (§5). To characterize these inconsistencies, we propose a taxonomy of five recurring annotation error categories grounded in discourse analysis: *fabricated relational links*, *scope and attachment errors*, *temporal and state change errors*, *hedging removal*, and *semantic reframing* (§6). We then analyze the downstream impact of these revisions on system evaluation (§7), finding that even small corrections yield clearer separation between competing evaluation approaches. This suggests that annotation inconsistencies can obscure meaningful differences in system performance.<sup>1</sup> Finally, we offer recommendations for future annotation efforts in long-form summarization, including the adoption of discourse-aware tools for quality assurance and improved design of error taxonomies and annotation frameworks (§8).

## 2 Related Work

**Annotation Reliability in Faithfulness Benchmarks.** Recent studies have raised concerns about the reliability of “gold-standard” annotations used in faithfulness evaluation benchmarks. [Laban et al. \(2023\)](#) conduct a manual analysis of model–dataset disagreements and estimate that at least 6% of samples in AggreFact ([Tang et al., 2023](#)), a widely used benchmark, are mislabeled. Similarly, [Seo et al. \(2025\)](#) report that in aggregated factuality datasets, 9.1% of examples are ambiguous and 6.6% appear mislabeled. Beyond annotation noise, [Godbole and Jia \(2025\)](#) show

<sup>1</sup>The revised datasets and codebase are available at <https://github.com/cs329yangzhong/faithfulEval4LongSumm>.

that state-of-the-art LLM-based faithfulness evaluators often rely on surface-level similarity, failing to properly verify claims that require aggregating information across distant parts of a source document. This limitation calls into question the reliability of methods that rely on LLM-generated rationales to guide human assessments, such as [Lee et al. \(2024\)](#). Taken together, these findings highlight the need for a more careful and systematic examination of annotation reliability in faithfulness benchmarks. *We address this gap through a discourse-aware reassessment of recent long-form summary datasets, supported by targeted human verification.*

### **Long-form Summary Annotation Granularity.**

As benchmarks expand to longer summaries, researchers have increasingly adopted fine-grained annotation schemes, including sentence-level error taxonomies ([Koh et al., 2022](#)) and span-level annotations with taxonomies designed to capture coherence-related errors in narrative texts ([Goyal et al., 2022](#)). [Krishna et al. \(2023\)](#) further show fine-grained clause-level annotations yield better inter-annotator agreement than holistic summary-level judgment. More recent long-form summary benchmarks have continued this trend by adopting sentence-level annotations, including both binary ([Subbiah et al., 2024](#); [Fang et al., 2024](#)) and fine-grained labels ([Ding et al., 2025](#)). Yet even fine-grained taxonomies operate primarily at the sentence level and rarely require annotators to verify relations between claims/sentences (e.g., causal or temporal). *Our work investigates this overlooked limitation by examining how human annotations fail to capture discourse-level inconsistencies.*

### **Discourse Analysis in Faithfulness Evaluation.**

[Pagnoni et al. \(2021\)](#) incorporate discourse link errors into their FRANK typology, marking an early effort to account for discourse-level phenomena in faithfulness evaluation. Building on this perspective, [Zhong and Litman \(2025a\)](#) show that faithfulness errors in long-form summaries correlate with specific discourse features such as nuclearity. Subsequent work has further incorporated discourse relations into structured faithfulness evaluation ([Zhang et al., 2025](#)). *In contrast, our work leverages a discourse-informed framework to produce fine-grained labels with interpretable, discourse-grounded rationales. We employ this framework as an auditing tool to surface systematic annotation errors that are often overlooked by standard evaluation procedures.*

	Text Diversity		Human Annotation			Dataset Statistics		
	Doc.Src	Taxonomy	Annotator	Protocol	Aggregation	# Summ	# Sent	Sum.Word
<b>STORYSUMM</b>	narratives	binary	6 crowd workers + 3 experts	hybrid human-AI	unanimity + adjudication	96	580	139
<b>VERIGRAY</b>	news	7-category	2 grad. students	double annot. + LLM review	disagreement review	412	2,044	122
<b>FAREBIO</b>	scientific	binary	2 medical doctors	single annotation	n/a	175	1,445	198

Table 1: **Overview of the three analyzed benchmarks.** We report the source document type (Doc.Src), annotation taxonomy, annotator type, protocol, and label aggregation strategy, along with key statistics: number of summaries (# Summ), number of annotated sentences (# Sent), and average summary length in words (Sum.Word).

### 3 Datasets

We analyze three faithfulness benchmarks for LLM-generated summaries which vary in domain, annotation protocol, and label taxonomy: STORYSUMM, VERIGRAY, and FAREBIO. Covering narrative, news, and scientific texts and applying distinct annotation strategies, they allow us to examine how annotation design affects labeling reliability. Table 1 summarizes their key characteristics.

**STORYSUMM** (Subbiah et al., 2024) contains 96 narrative summaries (580 annotated sentences). Each summary sentence is first labeled by three<sup>2</sup> crowd workers with binary faithfulness judgments, with explanations provided for unfaithful cases. The annotation interface includes an “N/A, commentary” option for sentences that interpret narrative themes rather than summarize story plots. To enhance label quality, the dataset adds (i) an *expert* setting with author adjudication at the summary level, and (ii) a *hybrid* setting where GPT-4 generates potential inconsistencies to guide a second round of crowd annotations. Final labels are derived via unanimity and adjudication, yielding a public version where some sentence labels diverge from the original unanimity votes. For our analysis, we use the raw per-annotator files, including individual votes and explanations.

**VERIGRAY** (Ding et al., 2025) contains 412 news article summaries (2,044), annotated by graduate students following a fine-grained taxonomy (Appendix A.1). The annotation protocol involves double annotation, followed by LLM-based reviews with prompts adopted from Seo et al. (2025), and a final annotator discussion to resolve disagreements. Here we analyze the initial dataset release (October 2025).

**FAREBIO** (Fang et al., 2024) contains 175 sum-

<sup>2</sup>This accounts for the annotator label construction in the original paper (Subbiah et al., 2024). The authors include two additional annotation settings (Stage 2 in §4), with a total of six crowd workers and three experts involved in the entire annotation process (Table 1).

maries (1,445 sentences) of biomedical research abstracts and introduction texts, annotated by two medical doctors for faithfulness and factual hallucination. Annotations include highlighting evidence spans and free-text explanations. Only a small subset of the data is doubly annotated (34 sentences, Cohen’s  $\kappa = 0.48$ ); the remaining sentences are singly annotated. Further details on annotation procedures and dataset construction of the three datasets are provided in Appendix A.

### 4 Reassessment Procedure

Our central hypothesis is that discourse-level inconsistencies (fabricated relations, temporal reordering) are overlooked by standard annotation procedures, which prioritize verifying individual facts over the relationships between them. To test this, we design a three-stage reassessment pipeline grounded in discourse analysis.

**Auditing Tool.** Our primary auditing tool is DESCRIBE,<sup>3</sup> a discourse-aware faithfulness evaluation framework grounded in Rhetorical Structure Theory (RST, Mann and Thompson 1988). Given a source document and a summary, DESCRIBE decomposes the summary into atomic claims while preserving discourse relations (e.g., causal, attribution, temporal, contrast), and performs structured inference to verify both claims *and* their intra- and inter-sentential relations against the source. This process yields fine-grained error labels with explicit rationales, identifying cases where facts are correct but the links between them are unsupported. Figure 2 (A) shows this for the Figure 1 example.

**Stage 1: Automated Flagging.** We apply DESCRIBE to evaluate all summary sentences across the three datasets. A sentence is flagged whenever DESCRIBE’s prediction disagrees with the gold label. We additionally apply two complementary

<sup>3</sup><https://anonymous.4open.science/r/Describe-2C68/>

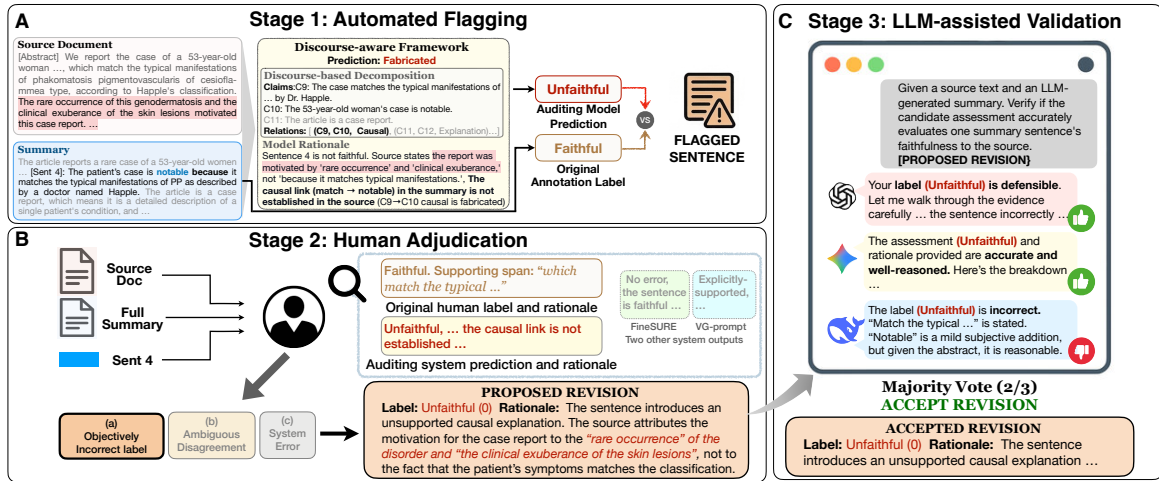


Figure 2: **Three-stage reassessment pipeline on a biomedical summary from FAREBIO.** Panel (A): a discourse-aware framework flags a fabricated causal link absent from the source. Panel (B): human adjudication reviews the source evidence and all system rationales and proposes a revised label with corresponding rationales. Panel (C): Three independent LLMs verify the proposed revision by majority vote. The example highlights a recurring pattern: individual facts are correct, but the relation connecting them is unsupported.

baselines to the flagged sentences to provide diverse rationales for Stage 2.

**Stage 2: Human Adjudication.** For each flagged instance, the first author reviewed the source document, full summary, gold label, annotator explanations (if available), and the auditing system’s predictions and rationales. To facilitate manual validation, we also apply two LLM-based faithfulness checkers with different prompting strategies—FineSURE (Song et al., 2024) and VG-prompt (Ding et al., 2025)—on those flagged sentences.<sup>4</sup> All systems use GPT-5 (medium reasoning effort) as the backbone. Each case is classified as: (a) *objectively incorrect*—the gold label contradicts guidelines or textual evidence; (b) *ambiguous disagreement*—both annotation and system rationale are defensible under reasonable interpretation, and (c) *system error*—the gold label is correct and the faithfulness checking systems err. For objectively incorrect cases, we propose revised labels aligned with the dataset taxonomy, with free-text rationales grounded in textual evidence, informed by DESCRIBE’s discourse-based rationales (see Figure 2 (B) for a worked example).

**Stage 3: LLM-assisted Verification.** To mitigate potential adjudication bias from a single annotator, all proposed revisions for “objectively incorrect” cases and their rationales are independently verified by three LLMs through web-based

chats: Gemini-3-Pro, GPT-5.2-Thinking, and DeepSeek-V3.2-Thinking.

As illustrated in Figure 2 (C), a revision is accepted only if it is supported by verifiable textual evidence (e.g., a misattribution, a numeric discrepancy, or an unsupported discourse relation) and corroborated by at least two models. This process ensures accepted changes reflect clear annotation inconsistencies rather than borderline disagreements.

## 5 Reassessment Outcomes

Table 2 presents detailed auditing results by category after applying our three-stage reassessment pipeline. We leave the ambiguous cases unchanged, focusing on clearly and objectively incorrect labels. We propose revised labels for 28 sentences in STORYSUMM (4.8% of 580), 69 in VERIGRAY (3.4% of 2,044), and 78 in FAREBIO (5.4% of 1,445). In STORYSUMM, where the raw per-annotator files are available, 15 of the 28 revised cases were flagged by at least one original crowd worker, indicating that valid minority signals were present but suppressed during label aggregation.

**Dataset-specific Patterns.** Revision patterns are not random but align with dataset-specific annotation protocols. In STORYSUMM, 21 of 28 revisions involve sentences originally labeled as faithful, reflecting the unanimity rule that labels a sentence as unfaithful only with full agreement—biasing the dataset toward under-detection of subtle errors. In VERIGRAY, most revisions occur at the

<sup>4</sup>Appendix B includes prompting details.

Dataset	Total Sent.	Flagged	Obj. Incorrect	Ambiguous	Inconsistency Rate
STORYSUMM	580	91	28	9	4.8%
VERIGRAY	2,044	98	69	4	3.4%
FAREBIO	1,445	115	78	5	5.4%

Table 2: Full auditing results after applying our three-stage reassessment pipeline presented in §4.

boundary between *Generally-Supported* and *Fabricated*, a distinction requiring judgment of when semantic strengthening becomes unsupported. In FAREBIO, where most sentences are annotated individually by medical doctors, revisions appear to reflect annotator-specific tendencies rather than aggregation effects. We observe repeated cases where surface-plausible comparative claims are accepted without verifying whether hedging in the source text has been preserved, suggesting that domain expertise may encourage acceptance of clinically consistent interpretations.

**Prevalence of Human Disagreement in STORYSUMM.** Using the official annotation files shared by the authors of STORYSUMM, we analyze the crowd-sourced annotation process and observe substantial disagreement among annotators. Specifically, 151 (out of 580) sentences received at least one label of “No” during annotation. Among these cases, the distribution of unfaithfulness votes—corresponding to cases where one, two, or all three annotators labeled a sentence as “No”—is 66.2%, 20.5%, and 13.2%, respectively. This suggests that unanimous agreement on unfaithfulness is relatively uncommon.

We further examine how these annotator votes are reflected in the aggregated sentence-level labels. For sentences where only one of three annotators assigned a “No” label, only 19% are ultimately labeled as “No” in the final dataset. In contrast, when two of three annotators labeled the sentence as “No”, the final label aligns with the majority judgment in 84% of cases. These findings indicate that the aggregation procedure substantially suppresses minority disagreement signals, leaving a large proportion of partially disagreed cases labeled as faithful. This observation suggests the presence of non-trivial annotation ambiguity and motivates targeted validation of these disputed instances.

## 6 Error Analysis

To characterize recurring patterns among the revised cases, we develop a taxonomy of five error types through qualitative analysis. This analysis

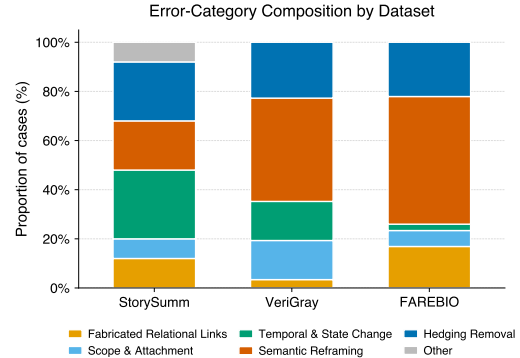


Figure 3: Normalized distribution of annotation error types across benchmarks.

identifies systematic sources of faithfulness errors that persist despite surface-level factual overlap. The resulting categories draw on established concepts from discourse processing and prior error taxonomies, including relational entailment, scope and attachment resolution, and temporal ordering (Mann and Thompson, 1988; Pagnoni et al., 2021).<sup>5</sup> Table 3 presents the five categories with diagnostic cues, discourse-based interpretations, and illustrative examples. Figure 3 shows their normalized distribution across datasets.<sup>6</sup> We highlight and discuss major observations across datasets below.

**Semantic reframing and hedging removal dominate across datasets.** In VERIGRAY and FAREBIO, most revisions involve either *semantic reframing*, which replaces factual descriptions with interpretive or evaluative language, or *hedging removal*, which omits epistemic qualifiers. Together, these categories account for roughly 70% of revisions. Because both preserve surface plausibility, they are difficult to detect when annotators judge whether content “sounds right” rather than whether the source framing is faithfully preserved.

We identify three subtypes of semantic reframing: (1) *interpretive labeling* (e.g., adding evalua-

<sup>5</sup>We used LLM-assisted summaries of the revised cases to surface candidate patterns, which were then manually consolidated and refined into the final taxonomy.

<sup>6</sup>Two additional STORYSUMM cases involve named entity errors outside the taxonomy and were corrected by the dataset authors after our notification.

Category	Typical cues	Definitions	Examples
<b>Fabricated Relational Links</b>	<i>because, due to, leading to, resulting in; but, instead of; if/unless</i>	Errors often involve <b>unsupported relations</b> (causal, contrastive, conditional) rather than missing facts.	<b>Summary:</b> “Casey leaves with the money <b>but</b> feels confused about the situation.” <b>Evidence:</b> Casey felt “confused” when asked to sign a name that wasn’t hers on her painting. After the argument she signed her own name, causing the requester to become “distracted” and “set about destroying the painting.” Casey then “said she was sorry and asked to be paid,” and “took the money and left.” <b>Error:</b> The word “but” fabricates a contrastive relation between the money and the confusion; the source presents them as separate, sequential events with no opposition.
<b>Scope &amp; Attachment</b>	Long multi-clause sentences; distant modifiers	Errors arise when connectives or modifiers attach to the <b>wrong proposition</b> , yielding false support under span matching.	<b>Summary:</b> Judge Ian Pearson remanded her into custody, stating she would likely be held at a male prison <b>due to</b> the risk of self-harm and non-attendance. <b>Evidence:</b> The judge said the defendant “would be a risk to herself and a risk of failing to attend,” thus ordering “remand in custody.” The reasoning for a male prison was that the defendant “had not made any physical changes or enhancements to her body or taken any medication.” <b>Error:</b> The summary incorrectly attaches the justification for the remand decision (self-harm risk, non-attendance) to the male-prison decision; each rationale belongs to a different proposition in the source.
<b>Temporal &amp; State Change</b>	<i>then/after/before; verbs like realize/discover/decide</i>	Errors occur when the correct fact is placed at the <b>wrong temporal position</b> in the narrative.	<b>Summary:</b> [Context:] When he hints at the gift, she gives him a blank stare and denies having ever done so. [Target:] They realize her memory is rapidly deteriorating. <b>Evidence:</b> The narrator recounted, “A few months ago I started to notice the memory problems”; the family shared the realization “we were so saddened about the Great Forgetting.” <b>Error:</b> The source places the narrator’s judgment about rapid deterioration <i>before</i> the hint exchange, not as a new realization triggered by it. The summary reverses the causal-temporal order.
<b>Hedging Removal</b>	<i>may aid vs are needed; mostly vs all; seemed vs is</i>	Errors arise when hedges or attributions are dropped, inflating <b>certainty beyond the source</b> .	<b>Summary:</b> The degradation correction <b>is</b> a manual process based on assumed knowledge of the sensor hardware. <b>Evidence:</b> “Degradation correction has been <b>mostly</b> a manual process based on assumed knowledge of the sensor hardware.” <b>Error:</b> The summary drops the hedge “mostly” and states the claim categorically, inflating certainty beyond what the source supports.
<b>Semantic Reframing</b>	interpretive wording (e.g., <i>high-risk, effective, promising</i> ); terminology reinterpretation; paraphrase drift	Errors arise when the summary <b>reframes a factual statement with an interpretive or evaluative label</b> not explicitly supported by the source.	<b>Summary:</b> It could also correctly identify artificially mutated viral genomes as <b>high-risk</b> . <b>Evidence:</b> “The artificial negative data with the replacement of the coding region of the spike protein were also predicted correctly (100% accuracy)” <b>Error:</b> By replacing the spike protein (the key to cross-species infection), the researchers created genomes that should not cause a pandemic. The model correctly identified these as negative (low risk),

Table 3: **Five recurring categories of overlooked annotation inconsistencies across datasets.** We present diagnostic cues, discourse-based detection insights, and illustrative examples with supporting source evidence.

tive terms not stated in the source like “high-risk”), (2) *domain-knowledge equivalence* (substituting terms based on unstated background knowledge), and (3) *interpretive inference* (drawing broader conclusions from observations). In FAREBIO, domain expertise appears to amplify this issue, as annota-

tors naturally accept clinically plausible interpretations without explicit textual support. Similarly, removing hedges (e.g., “mostly,” “felt like”) shifts qualified claims into a definitive claim—an episodic change that is easily overlooked despite preserved core content.

Backbone	Method	STORYSUMM				VERIGRAY				FAREBIO			
		Orig.	Rank	Rev.	Rank	Orig.	Rank	Rev.	Rank	Orig.	Rank	Rev.	Rank
Gemini-3-Flash	Fewshot	64.9	3	66.8	3	72.7	3	72.6	3	<b>70.9</b>	1	73.6	2 <sup>▽</sup>
	FineSURE	<b>70.1</b>	1	<b>69.8</b>	1	72.2	4	71.8	4	69.8	3	<b>74.2</b>	1 <sup>△</sup>
	VG-prompt	61.8	4	63.1	4	73.2	2	73.9	2	70.0	2	71.5	3 <sup>▽</sup>
	DESCRIBE	65.0	2	67.1	2	<b>75.5</b>	1	<b>76.8</b>	1	68.2	4	70.2	4
GPT-5	Fewshot	72.0	3	75.8	3	80.9	3	81.3	4 <sup>▽</sup>	<b>79.1</b>	1	85.5	3 <sup>▽</sup>
	FineSURE	<b>75.0</b>	1	77.2	2 <sup>▽</sup>	81.0	2	81.6	3 <sup>▽</sup>	77.2	3	82.7	4 <sup>▽</sup>
	VG-prompt	72.9	2	75.5	4 <sup>▽</sup>	80.6	4	82.5	2 <sup>△</sup>	78.5	2	85.7	2
	DESCRIBE	71.4	4	<b>77.5</b>	1 <sup>△</sup>	<b>81.5</b>	1	<b>84.9</b>	1	76.0	4	<b>87.3</b>	1 <sup>△</sup>

Table 4: **Balanced accuracy (%) and system rankings before (Orig.) vs. after (Rev.) label revision across three datasets and two backbone models.** Best results per model group are **bold**. <sup>△</sup> and <sup>▽</sup> mark methods whose rank improved or dropped after label revision. Predictions are unchanged; differences reflect corrected annotations. These shifts illustrate how annotation inconsistencies can obscure meaningful system rankings.

**STORYSUMM exhibits a distinct narrative-driven error profile.** Unlike the other datasets, revisions in STORYSUMM are more evenly distributed across categories, with *temporal and state-change* errors being most prevalent. This reflects the narrative text type: stories rely heavily on event order and the timing of character realizations, so misplacing an otherwise correct fact at the wrong narrative moment can distort its significance. This observation aligns with prior work emphasizing the importance of temporal discourse structure in narrative understanding (Genette, 1983; Hamilton et al., 2025). In contrast, temporal errors are rare in VERIGRAY and FAREBIO, where news and scientific texts are organized more thematically than chronologically, reducing opportunities for reordering errors. *Fabricated relational links* also occur relatively frequently, particularly in narrative summaries that compress complex plot structure into simplified causal or contrastive structures not present in the source document (e.g., the inserted “but” in the Casey example, Table 3, row 1).

**A common thread: correct facts, unsupported relations.** Across all categories and datasets, revised cases share a consistent pattern: individually correct facts paired with unsupported or distorted *relations* between them. For *fabricated relational links*, both cause and effect may appear in the source, but the causal connection is invented. Similarly, the introduction of contrastive relations (first example in Table 3), using discourse markers such as *instead of* or *but*, creates unsupported contrasts between otherwise valid facts. For *scope and attachment* errors, all facts are present, but modifiers attach to the wrong proposition. This pattern helps explain why such inconsistencies evade standard

annotation: verifying that individual facts appear in the source does not guarantee that the relationships between them, such as causal, temporal, attributive, or evaluative, are also supported.

Binary taxonomies (faithful/unfaithful) exhibit higher revision rates (4.8–5.4%) than VERIGRAY’s finer-grained seven-category scheme (3.4%), suggesting that additional distinctions help annotators recognize borderline cases. However, even fine-grained taxonomies remain vulnerable when applied at the sentence level without explicit verification of discourse relations. These observations motivate our recommendations in §8.

## 7 Annotation Revisions Enhance System Comparisons

**Experimental Setup.** We evaluate downstream effects by comparing system performance under original vs. revised labels for four faithfulness checking methods: Fewshot (Seo et al., 2025), FineSURE (Song et al., 2024), VG-prompt (Ding et al., 2025), and DESCRIBE using two representative backbone LLMs (Gemini-3-Flash and GPT-5). Performance is measured by sentence-level balanced accuracy, and we report system rankings per backbone and dataset. Experimental details are documented in Appendix C.

**Observations.** Table 4 shows that correcting a small number of annotations leads to meaningful shifts in system rankings. Of 24 system–dataset–backbone slots, 12 exhibit rank shifts after revision, with the effect concentrated under GPT-5 (9 of 12 slots) compared to Gemini-3-Flash (3 of 12). Under the original labels, system rankings are inconsistent: FineSURE (with GPT-5) ranks first on STORYSUMM, DESCRIBE on VERIGRAY,

and Fewshot on FAREBIO. After revision, rankings shift substantially. With **GPT-5**, DESCRIBE moves from rank 4 to rank 1 on STORYSUMM (71.4 → 77.5) and FAREBIO (76.0 → 87.3), while maintaining rank 1 on VERIGRAY. Most methods improve after revision, suggesting that corrected annotations yield more accurate and consistent evaluation, in line with Seo et al. (2025).

DESCRIBE benefits the most as the revisions primarily address discourse-level errors (e.g., fabricated relations, hedging removal, and misattached modifiers), which DESCRIBE is designed to detect. Under the original annotations, correct predictions on such cases were often penalized, suppressing its measured performance. Correcting these labels removes this penalty and reveals previously obscured differences, a pattern that holds across backbones but is most pronounced with GPT-5.

A potential concern is circularity, as DESCRIBE with GPT-5 was used to flag candidate cases. However, flagging only determined which cases were reviewed. Revised labels were established through human inspection of the source evidence and verified by three independent LLMs (Stages 2–3), rather than by simply adopting DESCRIBE’s predictions. Moreover, the revisions also do not uniformly benefit DESCRIBE (e.g., it remains fourth with Gemini-3-Flash on FAREBIO). These observations suggest that the revisions primarily correct discourse-level annotation errors rather than favoring a specific method.

## 8 Recommendations

Our findings align with prior work questioning the reliability of gold-standard annotations (Laban et al., 2023; Seo et al., 2025) and further identify a systematic pattern: in long-form summaries, annotation errors concentrate at the level of discourse relations rather than standalone individual facts. We distill four recommendations, ordered from post-hoc corrections to annotation design.

**Recommendation 1: Use discourse-aware tools for post-annotation quality assurance.** Our pipeline uncovers inconsistencies missed by both human annotators and existing evaluation systems. The key enabler is not model accuracy alone, but discourse-grounded rationales that make disagreements interpretable for human review. We therefore recommend applying discourse-aware evaluation tools as a post-annotation check, prioritizing manual review of cases with label disagreements. This

would be particularly valuable for *fabricated relational links* and *scope and attachment* errors (§6), where individual facts are present in the source but their relations are not supported.

**Recommendation 2: Reconsider unanimity-based aggregation.** In STORYSUMM, 21 of 28 confirmed revisions were false negatives—sentences incorrectly labeled as faithful—often flagged by a single annotator whose annotation was overridden by unanimity (§5). For discourse-rich summaries, where relational errors are inherently harder to detect, unanimity can suppress valid minority judgments. We recommend explicitly reviewing dissenting annotations to better balance precision and recall. More broadly, given the subjectivity inherent in genre-dependent datasets (Subbiah et al., 2025), datasets should retain annotator rationales and explore methods that model label variation (Plank, 2022; Weber-Genzel et al., 2024).

**Recommendation 3: Decompose claims before annotation.** Our error analysis shows that annotators reliably verify individual facts but often miss the relationships between them, as reflected in frequent *semantic reframing* and *hedging removal* errors (§6). Our human verification shows that presenting pre-decomposed atomic claims alongside their discourse relations can reduce annotators’ cognitive load and improve detection. This extends prior work on fine-grained annotation (Krishna et al., 2023) to a more linguistically grounded level by explicitly separating propositional content from epistemic and evaluative framing.

**Recommendation 4: Incorporate relation-level categories.** Many inconsistencies arise at the relation level—within sentences (misattached modifiers, dropped hedges) or across sentences (e.g., unsupported causal or temporal links)—yet all three benchmarks operate primarily at the sentence level, often with binary labels. While VERIGRAY’s finer taxonomy yields fewer revisions (3.4%), it still lacks explicit relation-level categories such as *fabricated causality* or *temporal reordering*. We recommend augmenting sentence-level labels with annotations of discourse relations (e.g., causal, temporal, attributive, evaluative) to capture whether these links are supported by the source.

## 9 Conclusion

We present a reassessment of three faithfulness benchmarks for long-form summaries, showing

that 3.4–5.4% of sentence-level annotations exhibit discourse-level inconsistencies. We introduce a five-category taxonomy—*fabricated relational links, scope and attachment errors, temporal re-ordering, hedging removal, and semantic reframing*—that explains how such errors arise and how a discourse-aware tool can detect them. Error patterns vary by genre and annotation design: narrative texts show more temporal and relational errors, while scientific and news texts are dominated by semantic reframing and hedging removal. Finer-grained taxonomies reduce but do not eliminate inconsistencies, and domain expertise can introduce biases through reliance on background knowledge. We propose four recommendations for future annotation efforts and release our revised labels for more reliable benchmark development.

## Limitations

Our findings are limited to English-language summaries. Generalization to other domains and languages can be a promising future direction. In addition, error analysis and categorization were primarily conducted by a single annotator. To mitigate potential subjectivity, revisions are grounded in explicit textual evidence and supported by structured, discourse-based rationales, enabling transparent verification. We further validated our STORYSUMM corrections against the original annotators’ free-text explanations, finding that 15 of 28 cases were independently flagged by at least one crowd worker. Finally, the auditing tool depends on discourse parsing quality, and errors in RST extraction may propagate to faithfulness judgments. More broadly, faithfulness evaluation remains inherently subjective at fine granularity, and revised human rationales may reflect biases of the underlying LLMs. To assess adjudication reliability directly, we had a second evaluator independently re-adjudicate a random sample of 15 flagged cases from the three datasets, reaching full agreement.

## Ethical Considerations

The datasets used in this paper are all publicly available for research purposes. We acknowledge the potential for bias in human evaluation.

**Responsible Disclosure.** We contacted the authors of all three datasets prior to publication. The authors of STORYSUMM (Subbiah et al., 2024) confirmed several issues and shared raw annotation files to support our analysis. We release our revised

labels alongside this paper to support transparent and reproducible benchmark development, rather than to undermine prior annotation efforts. The authors of VERIGRAY (Ding et al., 2025) updated the annotations based on our feedback, demonstrating the practical impact of our findings. The authors of FAREBIO (Fang et al., 2024) kindly provided the unprocessed annotation files but were unable to revise the dataset due to limited access to domain experts. Further validation with domain experts would strengthen revisions to FAREBIO and remains future work.

**Annotator Attribution.** We attribute observed errors to annotation protocol design (e.g., aggregation rules, taxonomy granularity) rather than individual annotators. We emphasize that the original annotators—crowd workers, graduate students, and medical doctors—operated within given guidelines; the identified issues reflect *procedural* limitations, not annotator competence.

**Intended Use.** Our revised labels and error taxonomy are intended to study the impacts of discourse-level errors in long-form summary evaluation and inform future annotation designs and practices. They are not meant to discredit existing benchmarks, as the vast majority of annotations (94.6–96.6%) remain unchanged after our reassessment.

## Acknowledgments

This research was supported in part by the University of Pittsburgh Center for Research Computing and Data, RRID:SCR\_022735, through the H2P cluster, which is supported by NSF award number OAC-2117681. We want to thank the members of the Pitt PETAL group and anonymous reviewers for their valuable suggestions, which helped improve this work. We also thank Zhehan Tiffany Zhu for her contributions to and discussions of the human evaluation.

## References

- Qiang Ding, Lvzhou Luo, Yixuan Cao, and Ping Luo. 2025. The gray zone of faithfulness: Taming ambiguity in unfaithfulness detection. *arXiv preprint arXiv:2510.21118*.
- Alexander R. Fabbri, Wojciech Kryściński, Bryan McCann, Caiming Xiong, Richard Socher, and Dragomir Radev. 2021. *SummEval: Re-evaluating summarization evaluation*. *Transactions of the Association for Computational Linguistics*, 9:391–409.

- Biaoyan Fang, Xiang Dai, and Sarvnaz Karimi. 2024. [Understanding faithfulness and reasoning of large language models on plain biomedical summaries](#). In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 9890–9911, Miami, Florida, USA. Association for Computational Linguistics.
- G rard Genette. 1983. *Narrative discourse : an essay in method / G rard Genette ; translated by Jane E. Lewin ; foreword by Jonathan Culler*. Cornell paperbacks. Cornell University Press, Ithaca, N.Y.
- Ameya Godbole and Robin Jia. 2025. [Verify with caution: The pitfalls of relying on imperfect factuality metrics](#). In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 22889–22912, Vienna, Austria. Association for Computational Linguistics.
- Tanya Goyal, Junyi Jessy Li, and Greg Durrett. 2022. [SNaC: Coherence error detection for narrative summarization](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 444–463, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Sil Hamilton, Matthew Wilkens, and Andrew Piper. 2025. [Narrabench: A comprehensive framework for narrative benchmarking](#). *Preprint*, arXiv:2510.09869.
- Huan Yee Koh, Jiaxin Ju, He Zhang, Ming Liu, and Shirui Pan. 2022. [How far are we from robust long abstractive summarization?](#) In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 2682–2698, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Kalpesh Krishna, Erin Bransom, Bailey Kuehl, Mohit Iyyer, Pradeep Dasigi, Arman Cohan, and Kyle Lo. 2023. [LongEval: Guidelines for human evaluation of faithfulness in long-form summarization](#). In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 1650–1669, Dubrovnik, Croatia. Association for Computational Linguistics.
- Wojciech Kryscinski, Nitish Shirish Keskar, Bryan McCann, Caiming Xiong, and Richard Socher. 2019. [Neural text summarization: A critical evaluation](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 540–551, Hong Kong, China. Association for Computational Linguistics.
- Philippe Laban, Wojciech Kryscinski, Divyansh Agarwal, Alexander Fabbri, Caiming Xiong, Shafiq Joty, and Chien-Sheng Wu. 2023. [SummEdits: Measuring LLM ability at factual reasoning through the lens of summarization](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 9662–9676, Singapore. Association for Computational Linguistics.
- Yuhoo Lee, Taewon Yun, Jason Cai, Hang Su, and Hwanjun Song. 2024. [UniSumEval: Towards unified, fine-grained, multi-dimensional summarization evaluation for LLMs](#). In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 3941–3960, Miami, Florida, USA. Association for Computational Linguistics.
- William C. Mann and Sandra A. Thompson. 1988. [Rhetorical structure theory: Toward a functional theory of text organization](#). *Text - Interdisciplinary Journal for the Study of Discourse*, 8(3):243–281.
- Artidoro Pagnoni, Vidhisha Balachandran, and Yulia Tsvetkov. 2021. [Understanding factuality in abstractive summarization with FRANK: A benchmark for factuality metrics](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4812–4829, Online. Association for Computational Linguistics.
- Barbara Plank. 2022. [The “problem” of human label variation: On ground truth in data, modeling and evaluation](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 10671–10682, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Wooseok Seo, Seungju Han, Jaehun Jung, Benjamin Newman, Seungwon Lim, Seungbeen Lee, Ximing Lu, Yejin Choi, and Youngjae Yu. 2025. [Verifying the verifiers: Unveiling pitfalls and potentials in fact verifiers](#). In *Second Conference on Language Modeling*.
- Hwanjun Song, Hang Su, Igor Shalyminov, Jason Cai, and Saab Mansour. 2024. [FineSurE: Fine-grained summarization evaluation using LLMs](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 906–922, Bangkok, Thailand. Association for Computational Linguistics.
- Melanie Subbiah, Faisal Ladhak, Akankshya Mishra, Griffin Thomas Adams, Lydia Chilton, and Kathleen McKeown. 2024. [STORYSUMM: Evaluating faithfulness in story summarization](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 9988–10005, Miami, Florida, USA. Association for Computational Linguistics.
- Melanie Subbiah, Akankshya Mishra, Grace Kim, Liyan Tang, Greg Durrett, and Kathleen McKeown. 2025. [Is the top still spinning? evaluating subjectivity in narrative understanding](#). In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 185–203, Suzhou, China. Association for Computational Linguistics.

- Liyan Tang, Tanya Goyal, Alex Fabbri, Philippe Laban, Jiacheng Xu, Semih Yavuz, Wojciech Kryscinski, Justin Rousseau, and Greg Durrett. 2023. [Understanding factual errors in summarization: Errors, summarizers, datasets, error detectors](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 11626–11644, Toronto, Canada. Association for Computational Linguistics.
- Leon Weber-Genzel, Siyao Peng, Marie-Catherine De Marneffe, and Barbara Plank. 2024. [VariErr NLI: Separating annotation error from human label variation](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2256–2269, Bangkok, Thailand. Association for Computational Linguistics.
- Kun Zhang, Oana Balalau, and Ioana Manolescu. 2025. [Structured discourse representation for factual consistency verification](#). In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 820–838, Vienna, Austria. Association for Computational Linguistics.
- Yang Zhong and Diane Litman. 2025a. [Discourse-driven evaluation: Unveiling factual inconsistency in long document summarization](#). In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 2050–2073, Albuquerque, New Mexico. Association for Computational Linguistics.
- Yang Zhong and Diane Litman. 2025b. [A tale of evaluating factual consistency: Case study on long document summarization evaluation](#). In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 12511–12532, Vienna, Austria. Association for Computational Linguistics.

## A Benchmark Details

We include additional details of the three studied benchmarks in the sections below.

### A.1 VERIGRAY

The error taxonomy of VERIGRAY is included in Table 5. Following the original paper’s suggestions, we exclude the ambiguous and “no-fact” cases in the experiments presented in §7, as these changes are hard to categorize and assess.

### A.2 STORYSUMM

The authors ask three crowd-sourcing annotators to assess whether each summary sentence is consistent with the original story. To reduce the influence of subjective “commentary” sentences (e.g., “*The story reflects the enduring bonds of friendship ...*” which interpret story themes rather than describe story plot), the annotation interface provides an additional “N/A, just commentary” label in addition to the “yes” or “no” options. The initial inter-annotator agreement reaches an almost perfect agreement of Fleiss’ kappa of 0.85.

Recognizing the challenges of evaluating narrative summaries, the authors introduce two additional procedures. In the *Expert* setting, three authors adjudicate summary-level labels. In the *Hybrid* setting, the authors first employ GPT-4 to generate potential inconsistencies between the source document and the summary, after which a new group of crowd workers reassesses the summaries using these suggestions. The final dataset merges unfaithful labels agreed upon by all three original annotators with adjudicated expert labels and labels produced through this hybrid human–AI process.

Despite these quality control steps, some aspects of the annotation scheme remain unclear, as the author only released a final version. In particular, crowd workers could assign an “N/A, Commentary” label to subjective sentences, but how these labels are incorporated into the final binary scheme is not documented. To examine this dataset, we obtained the raw per-annotator files from the original authors, including individual votes and free-text explanations, which allow us to analyze the potential impact of these ambiguous cases.

While conducting small-scale error analysis on the STORYSUMM predictions, we notice that several LLM-based approaches consistently reached different predictions compared to the sentence-level labels in the original dataset. These discrep-

ancies are confirmed to be either incorrect labels or subjectively annotated cases (where readers can have different interpretations given the vagueness of presentation), as reported in the same authors’ follow-up work (Subbiah et al., 2025).

We contacted the original authors of the STORYSUMM paper to validate our hypothesis that a certain level of annotation error exists in the original dataset. The authors confirmed that they collected three crowd-sourced annotator labels for each sentence and assigned the final label as “No” (unfaithful) only when all three annotators agreed; they also willingly shared the raw annotation files. We conduct the following analyses to validate the sentence-level labels, which have been overlooked in the original paper (they only apply diverse evaluation protocols on the summary-level).

**Ignoring Commentary Sentences.** Prior work analyzed STORYSUMM’s annotation files, finding that commentary sentences (at least one out of three annotators) are included in the released dataset. To align with the original paper’s definition, we exclude the same 19 sentences through from both summary-level and sentence-level assessments to maintain consistency with the evaluation guidelines of the original paper.

### A.3 FAREBIO

This dataset contains 175 LLM-generated summaries (1,445 sentences) of biomedical research abstracts and introductions, annotated by two medical doctors along two binary dimensions: faithfulness and factual hallucination, with highlighted supporting evidence spans. The reported inter-annotator agreement (Cohen’s Kappa = 0.48 on 34 double-annotated sentences) itself signals the difficulty of faithfulness judgment in this domain. While the binary scheme simplifies the annotation task, the limited disclosure regarding annotator training beyond the 34 example sentences and labeling guidelines, together with the reliance on single annotations for the majority of the dataset, leaves room for quality analysis.

We contacted the original authors and obtained a zip file containing the original annotation data. However, we were unable to recover the per-annotator annotation details due to the absence of documentation and the presence of complex annotation cache files.

Document	Target Sentence	Annotation
... qatar captial doha , home to the aspire dome , beat eugene to host the 2019 event in granting the championships to eugene the iaaf council ...	Eugene had previously failed in its bid to host the 2019 event, which was awarded to Doha.	Explicitly-Supported.
nicklaus holds up his ball to an adoring crowd as <b>gary player ( left ) and ben crenshaw salute the great crenshaw and nicklaus fist pump following his ace on the 130-yard</b>	<b>This feat was witnessed</b> by fellow golfers Gary Player and Ben Crenshaw.	Implicitly-Supported. <b>Reason:</b> The document supports the claim but does not explicitly mention the highlighted text.
A Gareth Anscombe <b>drop-goal edged Blues 23-20 ahead</b> after Gloucester Josh Hohnneck was yellow carded. But <b>unanswered second-half tries from Jonny May, Marshall, Mark Atkinson and Henry Purdy</b> sealed Gloucester’s win.	The Gloucester Rugby team <b>won 23-20 against Cardiff Blues</b> in a European Cup competition match.	Fabricated. <b>Reason:</b> The 23-20 is not the final score.
Italy’s National Institute of Geophysics and Volcanology (INGV) said <b>the quake struck at 15:48 (14:48 GMT)</b> , with its epicentre in Garfagnana.	The quake struck <b>at 15:48 GMT</b> and was followed by several aftershocks	Contradicting. <b>Reason:</b> It should be 14:48 GMT.
ormer Russian FSB colonel Igor Girkin, also known as Strelkov, who was then a key rebel commander in eastern Ukraine. Access to Anonymous International’s website is currently blocked in Russia.	3. <b>**Anonymous International (Shaltay Boltay)**</b> – Known for leaking Kremlin documents, the group has also released material on Ukraine, including emails from a <b>Russian rebel commander</b> .	Ambiguous. <b>Reason:</b> “russian rebel commander” is ambiguous: (russian rebel) commander vs. russian (rebel commander)
Francis I ( Franz Stefan , Francois Etienne 8 December 1708 – 18 August 1765 ) was <b>Holy Roman Emperor and Grand Duke of Tuscany , though his wife effectively executed the real powers of those positions</b>	The passage distinguishes between two historical figures named Francis I—one a <b>Habsburg emperor</b> and the other a Valois king of France—while focusing primarily on the latter’s reign, rivalries, and military strategies.”,	Out-Dependent <b>Reason:</b> Habsburg is not mentioned. However, <a href="https://en.wikipedia.org/wiki/Francis_I,_Holy_Roman_Emperor">https://en.wikipedia.org/wiki/Francis_I,_Holy_Roman_Emperor</a> writes that “Following the death of his father-in-law, Charles VI, in 1740, Francis and Maria Theresa became the rulers of the Habsburg domains.”
a pit crew member was hit by a car on sunday during the inaugural indycar grand prix of louisiana .	The incident occurred at the NOLA Motorsports Park in Avondale. <b>(Word count: 99) me know if you’d like any adjustments!</b>	No-Fact. <b>Reason:</b> This is considered as a meta note generated by LLMs, which is irrelevant to the summary and should be excluded from evaluation.

Table 5: **Examples of the taxonomy defined in VERIGRAY (Ding et al., 2025)**. The key segment that drives the annotation decision for each example is in red. We select these examples from the dataset.

## B Prompts

We include the prompts for FineSURE (Song et al., 2024) and a modified version of VG-prompt (Ding et al., 2025) in Table 6, which are adopted from the corresponding papers and codebases. For Fewshot, ClearCheck, and FaithLens, we used the model checkpoints open-sourced by the original authors and used their released prompts to process our data. For DESCRIBE, we use the publicly available tool.

## C Label Revision Experiments

**Backbones.** We picked two closed-source models, GPT-5<sup>7</sup> and Gemini-3-Flash<sup>8</sup>, and used their official APIs for our experiments. For GPT-5, we set the reasoning effort to medium due to budget constraints, while for Gemini-3-Flash, we used high reasoning effort. Due to computational cost, all results are averaged over two runs .

**Baselines.** We include the prompt for FineSURE (Song et al., 2024), a modified version of VG-prompt (Ding et al., 2025) in Table 6, adopted

<sup>7</sup>API name: gpt-5-2025-08-07

<sup>8</sup>API name: gemini-3-flash-preview

from the corresponding papers and codebases. For Fewshot, we follow the FEW\_SHOT\_TEMPLATE in the official codebase.<sup>9</sup> For DESCRIBE, we adopt the publicly available codebase.

**Datasets.** For the original datasets (prior to label revision), we experiment with the following versions:

- STORYSUMM: [https://github.com/melaniesubbiah/storysumm/blob/main/storysumm\\_w\\_subj.json](https://github.com/melaniesubbiah/storysumm/blob/main/storysumm_w_subj.json), retrieved in mid-February.
- VERIGRAY: <https://huggingface.co/datasets/Ding-Qiang/veri-gray-20251007>.
- FAREBIO: retrieved from <https://data.csiro.au/collection/csiro:63362>.

For the revised datasets:

- STORYSUMM: We correct the 28 reassessed labels and further exclude 19 “commentary

<sup>9</sup><https://github.com/justinseo/verifying-the-verifiers/blob/main/data/templates.py>

---

## FineSURE:

You will receive a transcript followed by a corresponding summary. Your task is to assess the factuality of each summary sentence across the following nine categories:

**no error:** the statement aligns explicitly with the content of the transcript and is factually consistent with it.

**out-of-context error:** the statement contains information not present in the transcript.

**entity error:** the primary arguments (or their attributes) of the predicate are wrong.

**predicate error:** the predicate in the summary statement is inconsistent with the transcript.

**circumstantial error:** the additional information (e.g., location or time) specifying the circumstance around a predicate is wrong.

**grammatical error:** the grammar of the sentence is so incorrect that it becomes meaningless.

**coreference error:** a pronoun or reference has a wrong or non-existing antecedent.

**linking error:** incorrect discourse linkage between statements (e.g., temporal or causal relations).

**other error:** the statement contains any factuality error not defined above.

### Instruction:

First, compare each summary sentence with the transcript.

Second, provide a single sentence explaining the factuality assessment.

Third, assign exactly one error category for each sentence.

Provide your answer in JSON format. The answer should be a list of dictionaries whose keys are "sentence", "reason", and "category":

```
[{"sentence": "first sentence", "reason": "your reason", "category": "no error"},  
{"sentence": "second sentence", "reason": "your reason", "category": "entity error"},  
{"sentence": "third sentence", "reason": "your reason", "category": "out-of-context error"}, ...]
```

### Transcript:

{DOCUMENT}

### Summary with N sentences:

1. {summary sentence 1 }

2. {summary sentence 2 }

...

N. {summary sentence N }

---

## VG-prompt

You are judging the faithfulness of each sentence of a summary to the source document. The faithfulness labels should be selected from the following options:

**A. Explicitly-Supported:** all atomic facts of the sentence appear verbatim (up to lexical or syntactic transformation) within the document.

**B. Generally-Supported:** the document entails the sentence, but the sentence is not explicitly supported. Minor differences are allowed only if the sentence adopts a weaker or less certain claim than the document. If any part of the sentence adopts a stronger or more certain claim, select Fabricated.

**C. Inconsistent:** the sentence logically contradicts the document.

**D. Fabricated:** the sentence does not contradict the document, but is neither implied by the document nor external world knowledge.

**E. Out-Dependent:** the sentence is not implied by the document alone, but is implied by combining the document with external world knowledge.

**F. Ambiguous:** the sentence or the document admits multiple interpretations.

**G. No-Fact:** the sentence does not contain factual content.

### Instruction:

Assess each summary sentence independently.

The output should be a JSON list enclosed within the special tag <FINAL\_PRED></FINAL\_PRED>. The list must contain one dictionary per sentence, with keys "sent\_id", "sentence", "rationale", and "faithfulness\_label".

Example output format:

```
<FINAL_PRED>  
[{"sent_id": 1, "sentence": "...", "rationale": "...", "faithfulness_label": "A"}, ...]  
</FINAL_PRED>
```

Ensure the number of output entries matches the number of summary sentences exactly.

### Document:

{DOCUMENT}

### Summary with N sentences:

1. {summary sentence 1 }

2. {summary sentence 2 }

...

N. {summary sentence N }

---

Table 6: FineSURE and VG-prompt prompts for faithfulness evaluation.

sentences” based on the analysis presented in Appendix A.2.

- VERIGRAY: Our findings were incorporated by the original authors, resulting in an updated release with additional author-side inspection. We therefore use the official updated version, 20251225.jsonl.
- FAREBIO: Because the authors were unable to review or validate the proposed revisions, we apply all our proposed corrections to the original dataset and use this modified version to reassess model performance.

**Evaluation Setup.** For VERIGRAY (Ding et al., 2025), we follow their first protocol to compute the balanced accuracy. Following their setup, we remove the Not Sure classes (Out-Dependent and Ambiguous) and No-Fact cases. The remaining classes can then be merged into two categories. Explicitly-Supported and Implicitly-Supported are merged into the **Faithful Class**, while Fabricated and Contradicting constitute the **Unfaithful Class**. We further report the balanced accuracy. Predictions are aligned accordingly: for methods with fine-grained predictions that follow the VERIGRAY taxonomy, specifically VG-prompt and DESCRIBE, classes with a faithfulness degree not less than a threshold (defined below) are considered **Faithful**, with the rest as **Unfaithful**. Here, the *order of faithfulness degree* is defined as:

$$\begin{aligned} & \text{Contradicting} < \text{Fabricated} < \text{Ambiguous} \\ & < \text{No-Fact} < \text{Out-Dependent} \\ & < \text{Implicitly-Supported} < \text{Explicitly-Supported} \end{aligned} \quad (1)$$

The threshold is set to the Implicitly-Supported class, following the original paper.

# Completing and Validating the Re-Aligned Switchboard Dialog Act Corpus

Run Chen<sup>1,2</sup>, Zihao Tao<sup>2</sup>, John Prado<sup>2,3</sup>, Ignazio LaManna<sup>2</sup>

Ryan Puterbaugh<sup>2</sup>, Mim Datta<sup>2</sup>, and Julia Hirschberg<sup>2</sup>

<sup>1</sup>Google, USA <sup>2</sup>Columbia University, USA <sup>3</sup>University of Alberta, Canada

julia@cs.columbia.edu

## Abstract

Although widely used in dialog act prediction and generation, the Switchboard Dialog Act (SwDA) corpus has performed poorly in models incorporating prosodic information because of misalignment between speech and text data. In this paper, we report our completion of the work begun in [Chen et al. \(2024\)](#) in addressing these misalignment issues with an improved SwDA corpus called RASwDA (Re-Aligned Switchboard Dialog Act Corpus). Now fully re-aligned and validated, RASwDA finally meets standards of accuracy allowing for classification models trained on it to exceed classification benchmarks set by models trained on other Switchboard subcorpora.

## 1 Introduction

Since its creation, the Switchboard Dialog Act (SwDA) corpus has been widely used for dialog act prediction and generation tasks. However, due to the misalignment between speech and text data in SwDA, models that incorporate prosodic information have shown poor performance. This paper presents the completed Re-Aligned Switchboard Dialog Act (RASwDA) corpus, as introduced in [Chen et al. \(2024\)](#), and reports novel state-of-the-art dialog act classification (DAC) results obtained directly from the efforts to manually correct and validate the previously force-aligned speech and text data. Over the DEVtest of [Chen et al. \(2024\)](#), this paper makes four contributions:

- We re-align the remaining 617.5 SwDA conversations and validate all 1,155, with a Cohen’s  $\kappa$  of 0.90 between the post-alignment and post-validation stages.
- We analyze inter-stage disagreement and find that Statement-Non-Opinion  $\leftrightarrow$  Statement-Opinion accounts for the majority of label changes.

- We train an acoustic-feature only BiLSTM model that saw a 41% improvement from training on the files from the original alignment to the post-alignment, showing that the incorrect alignments left prosodic features too noisy to learn from.
- We also train a late-fusion multi-modal RoBERTa model that sets a new DAC benchmark on SwDA (87.8% accuracy).

In §2, we outline some history of the Switchboard Telephone Speech Corpus. In §3, we summarize the procedure of [Chen et al. \(2024\)](#), present additional steps taken to validate the data set, and explain the models chosen for the DAC task. In §4 we present the results of an inter-annotator reliability test and classification models, and in §5 we discuss our conclusions.

## 2 Background

The **Switchboard Telephone Speech Corpus**, consisting of roughly 2,400 two-way telephone conversations between speakers of various US dialects, was created under DARPA sponsorship by Texas Instruments in 1990 but first released by the National Institute for Standards and Technologies in 1992 ([Godfrey et al., 1992](#)). The automated collection process prioritized the pairing of speakers who had not yet conversed either with each other or about the predetermined prompt. Participants initiated data-collection themselves, by giving a signal to begin the recording once prepared. The original corpus included time-aligned transcriptions of the speakers’ utterances, as well as the first attempt at automatic word-by-word time-stamping based on those transcriptions ([Wheatley et al., 1992](#)).

The original approach to aligning the transcriptions and DA tags with the appropriate utterances was a GMM-HMM Switchboard recognition system to generate the corresponding time intervals

(Shriberg et al., 1998, p. 454). However, consumer recording devices and automatic speech recognition systems have come a long way since then, and in fact the original outputs were so poorly aligned that they severely impaired research efforts in automatic tagging (Stolcke et al., 2000). Most force-alignment errors appear during quiet or low-energy utterances, but by manual inspection and verification of hundreds of audio files, we also find errors caused by background noise and line static. Similarly, the transcriptions contain incorrect or missing words, especially during speakers’ dialectical or nonstandard phrasings.

Since its creation, subcorpora of **Switchboard** have served as the basis for many re-annotations and re-segmentations. We discuss a selection of the most relevant to the present paper.

In 1997, researchers from UC Boulder and SRI International completed an annotation of roughly half of this corpus, LDC’s Switchboard-1 Release 2 (LDC97S62) (Jurafsky and Shriberg, 1997), with an augmented version of the Discourse Annotation and Markup System of Labeling (DAMSL) set of tags (Allen and Core, 1997). This **SwDA corpus** consists of 1,155 of the original conversations, and each of the roughly 205,000 utterances is tagged as one of 42 Dialog Acts, including Statement-opinion (sv), Acknowledge / Backchannel (b), Yes-No-Question (qy), etc.

By 1999, the third iteration of the Penn Treebank Project, **Treebank3**, included tagged, parsed, and disfluency-annotated text from 650 of the 2,400 Switchboard conversations (Marcus et al., 2000). This subcorpus formed the basis of the 2010 conversion of the corpus using the NITE XML Toolkit (hereafter **NXT**) (Carletta et al., 2005), which integrated different annotation protocols and allowed for easier traversal and querying of the data and its features (Calhoun et al., 2010). The benefits of this more complex representation included the ability to track long-distance dependencies in speech, to search the annotations as a set of attribute  $n$ -tuples, and to preserve the internal consistency of the constituent annotations.

In 2024, a team at Columbia University’s Department of Computer Science conducted a DEVtest to assess the efficacy of manual correction of previous forced alignments of DAMSL DA tags and corresponding Switchboard conversations (Chen et al., 2024). They used TextGrids, Praat’s tiered annotation format in which each tier holds a series

of labeled time-stamped intervals. This DEVtest included 537.5 of the 1,155 SwDA conversations, relying on **NXT-format Switchboard**’s XML files to create time-aligned TextGrids for the conversations included, and the *aneas* library for the rest (Chen et al., 2024).

### 3 Procedure

Our continuation of the work of Chen et al. (2024), to complete the **RASwDA** corpus, is the result of a year-and-a-half long process consisting of three stages: manual re-alignment of the remaining **SwDA** corpus, a manual validation of this re-alignment, and an evaluation stage. We follow the earlier data-preparation procedure (Chen et al., 2024), whose DEVtest re-aligned and validated 537.5 conversations (1075 individual speaker transcripts) of the 1,115 SwDA conversations (2,310 individual speaker transcripts): the remaining 617.5 conversations were drawn from both the **NXT-format Switchboard Corpus** (Calhoun et al., 2010) and LDC Switchboard-1 Release 2 (Godfrey et al., 1992). The **NXT-format Corpus** consists of 642 conversations in XML file-format, which we parsed into TextGrid format. For conversations not already in the **NXT Corpus**, we parsed each conversation’s transcript into separate transcripts, one for each speaker, and likewise separated each conversation’s audio into two WAV files using **SoX**. We computed the forced alignment for each utterance in each speaker transcript with the *aneas* library (Pettarin, 2017), which also resulted in a set of TextGrid files. The issues affecting forced alignment by *aneas* (e.g. background noises) observed in the DEVtest conducted by Chen et al. (2024) persisted into the remaining 617.5 conversations.

#### 3.1 Manual Re-alignment Stage

Like in Chen et al. (2024), we used the Praat speech analysis software (Boersma and Weenink, 2026) to manipulate these TextGrid-format SwDA transcripts. In addition to correcting the alignment of transcript time-intervals to the associated audio, the team of aligners was instructed to mark speaker overlap and laughter with special “SIL” and “<laughter>” tokens and to correct mis-transcriptions, other segmentation errors, and omissions in the transcript. To the best of our native-English-speaking judgment, we also resolved the mis-transcriptions and segmentation errors which the original SwDA annotators themselves marked

for correction at a later date (Jurafsky and Shriberg, 1997). Our aligners included six undergraduates, one post-baccalaureate, and one graduate student in computer science and linguistics programs, some compensated for their time in course credit.

### 3.2 Validation & Evaluation Stages

After completing manual re-alignment of the remaining 617.5 SwDA conversations, a validation stage began. A small team of validators (a subset of the aligners) reviewed all 1,155 conversations of the completed RASwDA corpus to verify that re-alignment had been performed for each conversation and to ensure any remaining marks were resolved. No student worked with the same file more than once from the re-alignment stage to the validation stage. General disagreements which validators brought against decisions made by aligners during the previous stage were discussed and addressed during team meetings. To formally assess annotator (dis)agreement among the aligners and validators, we computed the inter-stage reliability between post-alignment and post-validation DA tags. First, we extracted every DA tag from each TextGrid file in both versions of the corpus, recording all start times, end times, and labels. To compare the same utterance between the alignment and the validation stages, we matched tags by interval midpoint—rather than boundary—in order to tolerate segment boundary shifts of up to half the interval’s duration. The corresponding DA tags were then compared using Cohen’s Kappa, which measures agreement between the two annotation passes beyond what would be expected by chance.



Figure 1: Most-agreed-upon SwDA dialog acts between the alignment and validation stages.

To identify which DAMSL categories carry inherently opaque identification boundaries- information that downstream users of RASwDA need when interpreting model confusion on this corpus, we visualized inter-stage (dis)agreement and generated word bubbles for the DAMSL DA tags that were



Figure 2: Most-disagreed-upon SwDA dialog acts between the alignment and validation stages.

most- and least-frequently changed. For the agreement word bubble (Figure 1), we pulled the diagonal values from the confusion matrix (i.e., the labels that validators left unchanged), sorted by frequency, and kept the top 10 most-agreed-upon tags. We then mapped each DAMSL code to its full dialog act name using a dictionary built from the DAMSL manual in Jurafsky and Shriberg (1997) and generated a word bubble sized by frequency. For the disagreement word bubble (Figure 2), we pulled all off-diagonal values from the confusion matrix (i.e., where row  $\neq$  column: the labels that validators changed from the original to a new DA tag). We then summed each tag’s total outgoing changes and kept the top 10 most-disagreed-upon tags. Using the same DAMSL dictionary, we generated another word bubble sized by frequency.

### 3.3 Dialog Act Classification Models

As an extrinsic measurement of our realignment efforts, we compared the performance of two classification models, an acoustic only Bidirectional Long Short-Term Memory (Bi-LSTM), and a late fusion multi-modal RoBERTa model. We instantiate three instances of these models, trained on text grids from each stage of the re-alignment process. Both models utilize acoustic features extracted using the OpenSMILE’s low level descriptor (LLD) feature set, eGeMAPSv02.

#### 3.3.1 Bi-LSTM

The Bi-LSTM encoder processes variable-length acoustic feature sequences timestep-by-timestep, capturing context from both past and future frames simultaneously. We adopt a two-layer architecture with inter-layer dropout: the first layer learns low-level temporal patterns, while the second composes these into higher-level representations, with dropout regularizing the inter-layer activations. The encoder’s outputs are collapsed into a fixed-size ut-

terance embedding via learned attention pooling, which weights each frame by its relevance to the classification decision — down-weighting uninformative regions such as silence or background noise. Loss is computed using weighted cross-entropy, which inversely scales each class’s loss contribution to account for the severe label imbalance present in naturalistic dialogue data.

### 3.3.2 RoBERTa

The multi-modal model augments a pretrained RoBERTa-based text encoder with the same acoustic branch used in the Bi-LSTM model. The text encoder produces a 768-dimensional utterance representation from the [CLS] token, which aggregates sequence-level meaning during pretraining. The acoustic branch processes the eGeMAPSv02 feature sequences through an identical two-layer bidirectional LSTM with attention pooling, projecting the result to a 256-dimensional embedding. Rather than concatenating the two modalities directly, a learned scalar gate, conditioned on both the text and acoustic representations, controls how much the acoustic signal contributes for each utterance. This allows the model to down-weight noisy or uninformative acoustic signals when the text alone is sufficient, and to leverage prosodic cues when they are disambiguating. The gated acoustic embedding is concatenated with the text representation and passed through a two-layer MLP classifier trained on 10 dialogue act classes, with loss again computed using weighted cross-entropy.

## 4 Results

The Cohen’s Kappa of 0.90 falls within the ‘almost perfect’ range (0.81–1.00) on the Landis and Koch (1977) scale, reflecting strong inter-annotator reliability across post-alignment and post-validation annotations. The agreement bubble was dominated by Non-Verbal (12,232), Statement-Non-Opinion (5,456), and Backchannel (3,324). Despite ranking second in agreement, Statement-Non-Opinion also has the most outgoing changes in the disagreement bubble (723), with Statement-Non-Opinion to Statement-Opinion accounting for 578 of those changes, making it the primary source of annotator disagreement. That some DAMSL DA tags appear both frequently agreed-upon and frequently disagreed-upon we ascribe to a fuzziness inherent of the labels themselves.

## 4.1 Model Performance

Model	F1	Accuracy
<b>Bi-LSTM</b>		
Original	0.06924	15.536
Post-Align	0.30673	56.538
Post-Valid	0.29351	57.223
<b>RoBERTa</b>		
Original	0.72120	86.725
Post-Align	0.67939	87.600
Post-Valid	0.66812	<b>87.852</b>

Table 1: Model Results

As seen in Table 1, the results of our efforts produced significantly improved results. The Bi-LSTM is the more telling test of re-alignment because it relies on acoustic features alone. On the original data, it reached 15.5% accuracy, barely above chance for 42 classes. The original forced alignments mapped DA labels to the wrong portions of the audio, so the extracted features were effectively noise. After validation, accuracy rose to 57.2% over the three iterations of manual TextGrid correction, showing that corrected boundaries let the model pick up prosodic patterns it could not access before. The bold figure highlights a new state-of-the-art score for DAC models trained on the SwDA corpus, while the 41.7% increase in the performance from the original data to the post-validation data signifies that our model went from essentially utilizing random chance to a fully-functional classification model which relies on a number of temporal acoustic features from audio signals.

In addition, the performance distribution over the multi-modal DAC models shows that although text accounts for the most informative feature on classification for the accuracy of DAC tasks, corresponding acoustic features clearly provide a performance boost taking the model up from 86.7% to 87.8%.

## 5 Conclusion

This paper has presented the completed RASwDA corpus: the result of a large-scale manual re-alignment effort, as introduced in (Chen et al., 2024), to address issues present in SwDA (Jurafsky and Shriberg, 1997). With a Cohen’s  $\kappa$  of 0.90 and new DAC benchmarks on both acoustic-only and multi-modal models, we have shown the efficacy of manual alignment, annotation, and validation in

producing a well-aligned corpus. The validation stage confirmed the quality of the post-alignment data while also surfacing systematic label ambiguities, particularly between *sd* and *sv*, that users of RASwDA should be aware of.

## References

- James Allen and Mark Core. 1997. [Draft of DAMSL: Dialog act markup in several layers](#).
- Paul Boersma and David Weenink. 2026. [Praat: Doing phonetics by computer](#).
- Sasha Calhoun, Jean Carletta, Jason Brenier, Neil Mayo, Dan Jurafsky, Mark Steedman, and David Beaver. 2010. [The NXT-format Switchboard Corpus: A rich resource for investigating the syntax, semantics, pragmatics and prosody of dialogue](#). *Language Resources and Evaluation*, 44(4):387–419.
- Jean Carletta, Stefan Evert, Ulrich Heid, and Jonathan Kilgour. 2005. [The NITE XML Toolkit: Data model and query language](#). *Language Resources and Evaluation*, 39(4):313–334.
- Run Chen, Eleanor Lin, Shayan Hooshmand, Mariam Mustafa, Ritika Nandi Rose Sloan, Alicia Yang, Andrea Lopez, Ansh Kothary, Isaac Suh, Catherine Lyu, Eric Chen, Sophia Horng, and Julia Hirschberg. 2024. [RASwDA: Re-aligned switchboard dialog act corpus for dialog act prediction in conversation](#). In *[Proceedings] International Workshop on Spoken Dialogue Systems Technology*.
- John Godfrey, Edward Holliman, and Jane McDaniel. 1992. [SWITCHBOARD: Telephone speech corpus for research and development](#). In *[Proceedings] ICASSP-92: 1992 IEEE International Conference on Acoustics, Speech, and Signal Processing*, volume 1, pages 517–520.
- Dan Jurafsky and Elizabeth Shriberg. 1997. [Switchboard SWBD-DAMSL shallow-discourse-function annotation coders manual, draft 13](#).
- J. Richard Landis and Gary Koch. 1977. [The measurement of observer agreement for categorical data](#). *Biometrics*, 33(1):159–174.
- Mitchell Marcus, Beatrice Santorini, Mary Ann Marcinkiewicz, and Ann Taylor. 2000. [Treebank-3](#).
- Alberto Pettarin. 2017. [aeneas](#).
- Elizabeth Shriberg, Andreas Stolcke, Daniel Jurafsky, Noah Coccaro, Marie Meteer, Rebecca Bates, Paul Taylor, Klaus Ries, Rachel Martin, and Carol van Ess-Dykema. 1998. [Can prosody aid the automatic classification of dialog acts in conversational speech?](#) *Language and Speech*, 41(3–4):443–492. PMID: 10746366.
- SoX. 2015. [SoX: Sound eXchange](#).
- Andreas Stolcke, Klaus Ries, Noah Coccaro, Elizabeth Shriberg, Rebecca Bates, Daniel Jurafsky, Paul Taylor, Rachel Martin, Carol Van Ess-Dykema, and Marie Meteer. 2000. [Dialogue act modeling for automatic tagging and recognition of conversational speech](#).
- Barbara Wheatley, George Doddington, Charles Hemphill, John Godfrey, Edward Holliman, Jane McDaniel, and Drew Fisher. 1992. [Robust automatic time alignment of orthographic transcriptions with unconstrained speech](#). In *[Proceedings] ICASSP-92: 1992 IEEE International Conference on Acoustics, Speech, and Signal Processing*, volume 1, pages 533–536.

# Not Worth Mentioning? A Pilot Study on Salient Proposition Annotation

Amir Zeldes, Katherine Conhaim, Lauren Levine

Department of Linguistics

Georgetown University

{amir.zeldes, kc1512, le176}@georgetown.edu

## Abstract

Despite a long tradition of work on extractive summarization, which by nature aims to recover the most important propositions in a text, little work has been done on operationalizing graded proposition salience in naturally occurring data. In this paper, we adopt graded summarization-based salience as a metric from previous work on Salient Entity Extraction (SEE) and adapt it to quantify proposition salience. We define the annotation task, apply it to a small multi-genre dataset, evaluate agreement and carry out a preliminary study of the relationship between our metric and notions of discourse unit centrality in discourse parsing following Rhetorical Structure Theory (RST).

## 1 Introduction

Judging how important a proposition is to the contents of a text or conversation is important for a range of tasks in information retrieval, summarization and text mining (Narayan et al., 2018; Nguyen and Le, 2024; Dwivedi et al., 2025). In extractive summarization in particular, it is implied that the most important propositions in a document will be selected for inclusion in summaries. However, little work has been done to explain which propositions should be considered salient, and what the operationalization of salience should be, or how it could be measured.

By contrast, salience annotation for entities has been studied more extensively. While some papers have shown that manual annotation of salient entities exhibits low agreement (Dojchinovski et al., 2016; Trani et al., 2018), an alternative paradigm leveraging document summaries (Dunietz and Gillick, 2014) has been more successful in delivering a consistent operationalization by equating salience with **summary worthiness**: since it is difficult to summarize a document without mentioning its most salient entities, entities mentioned in a summary can be considered salient.

Two main criticisms of this approach have been recently addressed by Lin and Zeldes (2025): 1. that summaries themselves are subjective, meaning different summaries will result in different salience annotations; and 2. that the resulting annotations are binary (mentioned/not-mentioned in the summary). Lin and Zeldes address this by obtaining five summaries for each document, which allows for a gradient metric (entities mentioned in 5 summaries are more salient than those mentioned in 2 or 3), and controls for summary subjectivity by obtaining multiple judgments.

In this paper, we aim to test the same methodology to the concept of proposition salience. Specifically, we will define a gradient salience metric for propositions based on **how many summaries mention a proposition**. With the data at hand, we will provide some first insights on characteristics of salient propositions, and in particular, determine the extent to which the metric correlates with an existing notion of discourse unit centrality in the tradition of RST discourse parsing. The main contributions of this paper are:

- Definition and evaluation of a summary-based approach to graded proposition salience
- A pilot dataset of salient propositions aligned to summaries with additional annotations
- Analysis of the correlation between our approach and RST centrality
- An annotation interface for the task

## 2 Previous work

Much of the work on ranking proposition salience has been carried out in the context of summarization tasks (Saggion and Poibeau, 2013), specifically optimizing for selecting sentences that maximize NLG evaluation metrics (Nallapati et al., 2017; Narayan et al., 2018), such as ROUGE scores (Lin,

2004). Early work on sentence or paragraph importance in texts focused on direct human judgment and recall tasks testing the memorability of parts of a text in closed book scenarios (Trabasso and Sperry, 1985; Wolf and Gibson, 2004). Subsequent work on extractive summarization focused on identifying the most relevant sentences for direct inclusion in summaries. Seminal work by Nenkova and Passonneau (2004) showed that propositions included in multiple reference summaries should be prioritized and used in the evaluation of system summaries. Later studies shifted focus to supervised approaches including tree-based methods leveraging semantic annotations (Fang et al., 2016), end-to-end representation learning in embedding spaces (Chen et al., 2018), and more recently using LLM prompting (Parmar et al., 2024).

Studies using annotated data have relied either on direct annotation of important sentences (Liu et al., 2018), or on their derivation from extractive summarization datasets for sentence selection (Cheng and Lapata, 2016), both of which suffer from instability. Some previous work, such as (Liu et al., 2018), also attempted to use lexical overlap (lemmas) between documents and summaries to approximate binary salience for sentences, but did not carry out extensive manual evaluation or attempt to use multiple summaries. The most similar previous work to the present paper is Lin and Zeldes (2024, 2025), which uses a summary-based methodology to rank salient entities based on the number of summaries they appear in, including manual verification of alignments. Below we apply analogous methods to ranking salient propositions.

### 3 Data

The data used by Lin and Zeldes (2025) comes from the freely available GUM corpus (Zeldes, 2017), which includes five summaries per document. These were collected in an earlier project aiming to produce consistent and faithful summaries across a broad range of genres using controlled guidelines, which instructed summarizers to remain as close as possible to the text (Liu and Zeldes, 2023), thereby facilitating our task. We re-use that data for this pilot study, annotating the test set, which comprises 32 documents from 16 genres of spoken and written English with ~30K tokens and 3,800 propositions (see Table 1). Coupled with 5 summaries per document (and therefore per proposition), our data consists of close to 19K

proposition annotations (alignments).

Genre	Docs	Tokens	EDUs	Alignments
<i>academic</i>	2	1,952	250	1,250
<i>biography</i>	2	1,679	182	910
<i>conversation</i>	2	1,868	342	1,710
<i>court</i>	2	2,075	253	1,265
<i>essay</i>	2	2,359	301	1,505
<i>fiction</i>	2	2,029	264	1,320
<i>how-to</i>	2	1,642	223	1,115
<i>interview</i>	2	1,653	209	1,045
<i>letter</i>	2	1,939	224	1,120
<i>news</i>	2	1,891	200	1,000
<i>podcast</i>	2	2,119	257	1,285
<i>reddit</i>	2	1,858	257	1,285
<i>speech</i>	2	1,728	184	920
<i>textbook</i>	2	2,072	255	1,275
<i>travel</i>	2	1,722	154	770
<i>vlog</i>	2	1,669	224	1,120
<b>Total</b>	32	30,255	3,779	18,895

Table 1: GUM V12 test set data.

Additionally, the original corpus contain discourse trees following enhanced Rhetorical Structure Theory (eRST, see Mann and Thompson 1988 and Zeldes et al. 2025), which we use below to identify proposition boundaries. Since RST trees can be used to assess the centrality of discourse units based on their nesting depth relative to the tree root, we can evaluate and compare our salience annotations based on the nesting depth data from the trees (see Appendix B for details).

### 4 Annotation

Given a text and summaries, our annotation task involves three steps: identifying propositions, aligning them to summaries, and categorizing properties of the alignment.

**Defining proposition markables** Although propositional structure can be complex and potentially nested, our goal is to be able to rank spans of text for salience with unambiguous scores, meaning we require a ‘tiling’ approach, in which every word in the document belongs to exactly one proposition. To do this in a principled way, we rely on RST’s existing notion of Elementary Discourse Units (EDUs). At a maximum, an EDU can be no larger than a sentence, as in (1); this applies even if the sentence is a fragment, in which case we may need to treat a single noun phrase (for example a heading) in the same way as a ‘proposition’, as

in example (2). However we stress that this is inevitable, since summary content can be triggered by verbless utterances such as headings, which can then reasonably be seen as salient spans. Aside from such fragments, EDUs generally correspond to single-clause propositions, usually with one verb and its arguments (3). Additionally, EDUs may be discontinuous, for example due to interruptions by relative clauses (4). We use the pre-existing EDU segmentation in the GUM corpus without modifications as markable propositions, and treat discontinuous EDUs as single units.

- (1) [Kim called yesterday.]<sub>1</sub>
- (2) [INTRODUCTION]<sub>1</sub>
- (3) [They called]<sub>1</sub> [because you forgot it]<sub>2</sub>
- (4) [The boy]<sub>1</sub> [who laughed]<sub>2</sub> [fell asleep]<sub>1</sub>

We also considered the alternative of focusing on entire sentences, but realized these would be too coarse grained: sentences in the corpus contain an average of 2.28 EDUs, and our annotations below indicate that while most sentences are not included in summaries at all (about 78.5%), equal amounts of sentences are mentioned in their entirety (10.7%) versus partially (10.8%, in which some EDUs are mentioned and others are not), meaning sub-sentential alignment is common.

**Alignment** To decide whether a proposition appears in a summary, we rely on two levels of alignment: a strong alignment match is present if the summary refers to the same event or predicate in the proposition, or for a fragment, refers to the same entity. However in many cases, a proposition is not mentioned exactly, but it is essential for the content of the summary. For example, the summary might mention when the events in a document take place, but not mention the predicate that actually contained the time information. In such cases, we say that the proposition **triggers** content in the summary, and refer to the alignment as **approximate**. Our complete guidelines, which give a range of such examples, are made publicly available with this paper.<sup>1</sup>

**Categorization** Proposition alignment proceeds summary-wise, with each document-summary pair receiving a set of binary alignment judgments for each proposition. Once an alignment has been

made, we further categorize it into a proposition **match** versus **approximate** alignment. For approximate cases, we further sub-categorize alignments into normal triggers, where some information from the proposition is carried into the summary directly, and **component** alignments, in which a summary mentions an aggregation over propositions in the document. For example, if the summary of an article about where NASA shuttles will be displayed mentions ‘some people were disappointed with the final location decisions’, but no such statement appears in the document directly, the examples in (5) are taken to be components instantiating the triggering information in aggregate.

- (5) a. *While the museum of flight was in the top running, I’m disappointed that NASA did not choose them*
- b. *Cornyn’s statement added ... I’m deeply disappointed with the Administration’s misguided decision*

Neither of these two propositions appear in the summary directly, but they are the propositions which, taken together, result in the summary’s statement about some people’s disappointment.

Additionally, if multiple propositions can be made responsible for the same information in the summary (match or approximate), we annotate them as **duplicates**, indicating that while they provide salient information, there is substantial redundancy between them. Component alignments containing distinct information, as in example (5) are not considered duplicates.

**Annotation interface** To align propositions we create a dedicated interface called **GlowLyter**<sup>2</sup> which allows highlighting pre-defined units by clicking, with separate highlighting storage saved for each of the summaries displayed at the top (Figure 1). A note button next to each span opens the additional annotation box to indicate approximate, component and duplicate alignment information, which receive unique colors (yellow default highlighting, cyan for ‘approximate’ and green for ‘component’). The interface can easily be used to produce any kind of highlighting annotations of sentences or other spans and uses a file format that allows definition of custom annotation fields with text boxes or check boxes, making it adaptable to other types of annotations (see Appendix A). We

<sup>1</sup><https://gucorpling.org/amir/pdf/proposal-guidelines.pdf>.

<sup>2</sup><https://github.com/gucorpling/glowlyter>

release the code for the interface with our data on our GitHub repository under the open Apache 2.0 license. The annotated data will be released as part of the publicly available GUM corpus.

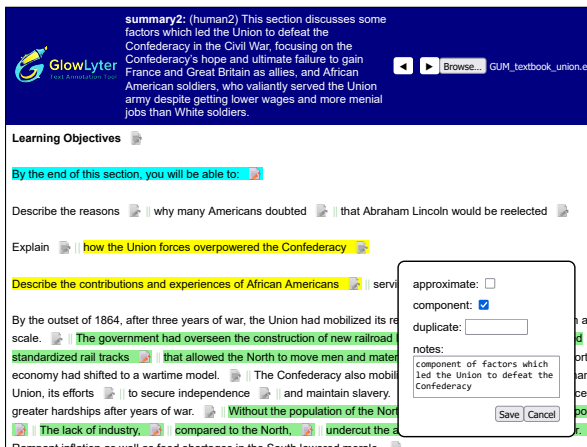


Figure 1: GlowLyter annotation interface for summary-wise salient proposition alignment.

## 5 Evaluation

Because most propositions are not salient ( $\sim 90\%$ , depending on the genre/document), we anticipate that % agreement (propositions two annotators assign the same label) is an optimistic metric; we therefore also report Cohen’s kappa, which takes into account the probability of chance agreement. We measure agreement in four increasingly lenient scenarios: 1. exact agreement, incl. all EDUs; 2. strictly literal agreement, disregarding disagreements involving components; 3. strictly matching agreement, the same but disregarding approximates; and 4. strictly matching duplicate-agnostic agreement, the same as 3. but treating ‘duplicate’ units as interchangeable (annotators agree if they both flag one in a set as salient, or both flag none).

metric	accuracy		$\kappa$	
	micro	macro	micro	macro
<i>strict all</i>	92.97	92.73	65.43	64.62
<i>strict literal</i>	95.08	94.78	74.07	72.74
<i>strict match</i>	95.96	95.73	78.05	77.57
<i>duplicates OK</i>	96.63	96.52	83.99	82.64

Table 2: Inter-annotator agreement

Table 2 shows that strict agreement is only moderate. Tolerating interchangeable duplicates (ways of mentioning ‘the same thing’) substantially raises agreement (+5.9% micro  $\kappa$ ). Component alignments are responsible for the most disagreements

(+8.6% micro  $\kappa$  when disregarded), much more than non-component approximates add (+3.98%).

Qualitatively we note cases where annotators focus more on lexical overlap than identical reference. For example for a summary mentioning ‘*a memory of the last time they saw a "yellow man"*’, one annotator selected ‘*I remembered the only "yellow man" I had ever seen*’, while another omitted ‘*I had ever seen*’, which references ‘seeing’, but may mean the specific seeing event or experience in general. Fragment component mentions also posed difficulties: one summary mentions a ‘family’, but the text enumerates only the members including two EDUs for the mother: ‘*my mother twisting her hands*’ and ‘*my father’s wife with red eyes*’. Different annotators preferred each of these: one argued for the first based on text order, while the other preferred the second since it did not include a superfluous verb ‘twisting’, a proposition that did not appear in the summary.

## 6 Analysis

Figure 2 shows the salience score distribution. Unlike summary-based entity salience scores, which follow a U-shaped distribution (Lin and Zeldes, 2025), proposition scores follow a descending histogram. Whereas, based on previous work, most entities are not in summaries, but the next biggest group is entities appearing in all summaries, the largest amount of salient propositions are mentioned in only one summary. This suggests inter-summary variance is higher at the level of proposition selection than is the case for participants.

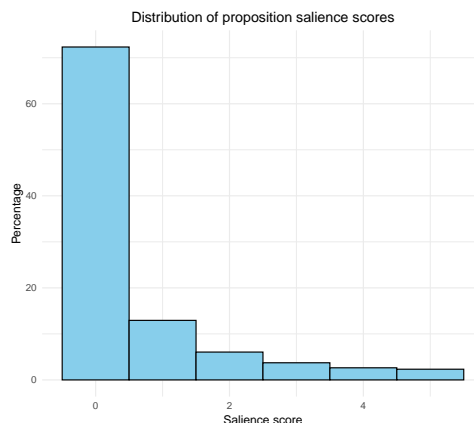


Figure 2: Distribution of salience scores.

Since we have RST trees for our data, we can also examine how our scores correspond to centrality in RST, i.e. how far an EDU is from the root

of the tree. We operationalize this distance as a proportion from 0 (the unit is the root) to 1 (the unit has the maximum distance in its document; see Appendix B and Figure 4 for details).

We find that RST centrality is significantly correlated with salience ( $r^2 = -0.287, p < 0.0001$ ) though the correlation is diminished by many confounds: longer EDUs are more likely to overlap with summaries, as are complete sentences compared to clauses. RST relations, which indicate the discourse role of each unit (expressing PURPOSE, CAUSE etc., see Appendix C) are also likely to influence salience, as does the EDU’s position in the document: the earlier a unit appears, the higher its score ( $r^2 = -0.119, p < 0.0001$ ).

In order to test whether RST centrality remains a significant predictor of proposition salience when these confounds are accounted for, we construct a mixed effects model with the document as a random effect. Since our salience scores do not follow a normal distribution (see Figure 2) and only allow for discrete integer values, we must use a target distribution that is compatible with non-normal, and specifically with long-tailed, integer data.<sup>3</sup> To that end, we use beta-binomial regression, modeling each summary containing a proposition as a single ‘success’, and each proposition as instantiating five trials (chances to be mentioned in a summary). To test the significance of individual features, we build the complete model with all features, and then perform single term deletions, comparing the full model to a model with each feature removed, using a Likelihood Ratio Test (LRT), reported in Table 3).

Feature	df	AIC	LRT	Pr( $\chi^2$ )
<none>		-4000.5		
<i>position</i>	1	-3962.8	39.68	2.987e-10 ***
<i>is_sent</i>	1	-3935.5	67.02	2.682e-16 ***
<i>length</i>	1	-3906.5	95.98	< 2.2e-16 ***
<i>relation</i>	31	-3469.0	593.46	< 2.2e-16 ***
<i>centrality</i>	1	-3198.2	804.28	< 2.2e-16 ***

Table 3: LRT single term deletions for features in a mixed effects beta-binomial regression model predicting salience, ranked by AIC.

The model confirms that centrality is the strongest predictor of salience when confounds are accounted for (maximum difference using Akaike’s

<sup>3</sup>We thank an anonymous reviewer for commenting on this.

Information Criterion, AIC). It is followed by the relation type and, with a much weaker effect, by unit length. However the model still achieves only a weak fit to the data (marginal  $r^2 = 0.131$ , conditional  $r^2 = 0.159$ ) and a binary classification accuracy of just 73.87% over a majority baseline of 72.33% (always predicting salience=0). This suggests much work remains to be done in understanding what characterizes salient propositions.

## 7 Conclusion and outlook

This paper presented a first effort to manually annotate gradient proposition salience based on alignment of discourse units with multiple summaries, adapting a methodology already in use for entities to the realm of predicates. Our evaluation shows that while the task is challenging, agreement is far above chance, and we plan to apply lessons from our adjudication process to refining our guidelines.

Work on salient entities (see Zeldes and Lin 2026) has revealed a range of expected and surprising results about what makes some parts of a text or conversation particularly noteworthy, but much corresponding work on propositions is still pending. The results of the preliminary study above already suggest a range of correlations between the notion of proposition salience developed in this paper and other directly observable properties (document position, length) and theoretical constructs (discourse relations, RST centrality) applying to propositions. Very recent concurrent work on entity salience has also shown correlations with the ways in which text is represented by language models, such as effects related to surprisal (Lin and Zeldes, 2026), and we are excited by the prospect of discovering additional properties that make propositions more or less salient in a similar vein.

We also intend to annotate the rest of the GUM corpus data and its complementary out-of-domain test set with challenging genres, GENTLE (GENre Tests for Linguistic Evaluation, Aoyama et al. 2023). We plan to apply further manual annotation to the dev set and use NLP/LLM-assisted annotation for the train set. The test set created here will be used to evaluate models for the task and develop prompting approaches for LLM ensembles to approach human performance. We hope that this data will enable detailed analyses of what makes propositions salient, and that our publicly released annotation guidelines and interface will be useful for the community.

## Limitations

This pilot study covers only the test set of one corpus, limited to contemporary English. Although we believe the quality of both the EDU segmentation and the summaries in GUM is high, our annotations are by definition only as good as these underlying sources of information. The error analysis in this short paper is limited to a few examples due to space. We also recognize that agreement on the proposition alignment task is still far from perfect, and are using our experiences from this study to refine guidelines for the final adjudicated data release. We look forward to feedback from the community on ways to make the data as consistent and reproducible as possible.

## References

- Tatsuya Aoyama, Shabnam Behzad, Luke Gessler, Lauren Levine, Jessica Lin, Yang Janet Liu, Siyao Peng, Yilun Zhu, and Amir Zeldes. 2023. **GENTLE: A genre-diverse multilayer challenge set for English NLP and linguistic evaluation**. In *Proceedings of the 17th Linguistic Annotation Workshop (LAW-XVII)*, pages 166–178, Toronto, Canada. Association for Computational Linguistics.
- Kuan-Yu Chen, Shih-Hung Liu, Berlin Chen, and Hsin-Min Wang. 2018. **An information distillation framework for extractive summarization**. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 26(1):161–170.
- Jianpeng Cheng and Mirella Lapata. 2016. **Neural summarization by extracting sentences and words**. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 484–494, Berlin, Germany. Association for Computational Linguistics.
- Milan Dojchinovski, Dinesh Reddy, Tomáš Klietr, Tomáš Vitvar, and Harald Sack. 2016. **Crowdsourced corpus with entity salience annotations**. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC’16)*, pages 3307–3311, Portorož, Slovenia. European Language Resources Association (ELRA).
- Jesse Dunietz and Daniel Gillick. 2014. **A new entity salience task with millions of training examples**. In *Proceedings of the 14th Conference of the European Chapter of the Association for Computational Linguistics, volume 2: Short Papers*, pages 205–209, Gothenburg, Sweden. Association for Computational Linguistics.
- Pulkit Dwivedi, Preet Kamal, Mansi Kajal, Nikita Rattan, and Namit Chawla. 2025. **Nyx: An extractive text summarization framework using TF-ISF and graph-based sentence ranking**. In *2025 13th International Conference on Intelligent Embedded, Micro-Electronics, Communication and Optical Networks (IEMECON)*, pages 1–5.
- Yimai Fang, Haoyue Zhu, Ewa Muszyńska, Alexander Kuhnle, and Simone Teufel. 2016. **A proposition-based abstractive summariser**. In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pages 567–578, Osaka, Japan. The COLING 2016 Organizing Committee.
- Chin-Yew Lin. 2004. **ROUGE: A package for automatic evaluation of summaries**. In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.
- Jessica Lin and Amir Zeldes. 2024. **GUMsley: Evaluating entity salience in summarization for 12 English genres**. In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2575–2588, St. Julian’s, Malta. Association for Computational Linguistics.
- Jessica Lin and Amir Zeldes. 2025. **GUM-SAGE: A novel dataset and approach for graded entity salience prediction**. In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 438–455, Vienna, Austria. Association for Computational Linguistics.
- Jessica Lin and Amir Zeldes. 2026. **Expect the unexpected? Testing the surprisal of salient entities**. In *Proceedings of the 64th Annual Meeting of the Association for Computational Linguistics (ACL 2026)*, San Diego, CA. Association for Computational Linguistics.
- Yang Janet Liu and Amir Zeldes. 2023. **GUMSum: Multi-genre data and evaluation for English abstractive summarization**. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 9315–9327, Toronto, Canada. Association for Computational Linguistics.
- Zhengzhong Liu, Chenyan Xiong, Teruko Mitamura, and Eduard Hovy. 2018. **Automatic event salience identification**. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1226–1236, Brussels, Belgium. Association for Computational Linguistics.
- William C. Mann and Sandra A. Thompson. 1988. **Rhetorical Structure Theory: Toward a functional theory of text organization**. *Text*, 8(3):243–281.
- Ramesh Nallapati, Feifei Zhai, and Bowen Zhou. 2017. **SummaRuNNer: A recurrent neural network based sequence model for extractive summarization of documents**. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence (AAAI-17)*, pages 3075–3081, San Francisco, CA.

- Shashi Narayan, Shay B. Cohen, and Mirella Lapata. 2018. [Ranking sentences for extractive summarization with reinforcement learning](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1747–1759, New Orleans, Louisiana. Association for Computational Linguistics.
- Ani Nenkova and Rebecca Passonneau. 2004. [Evaluating content selection in summarization: The pyramid method](#). In *Proceedings of the Human Language Technology Conference of the North American Chapter of the Association for Computational Linguistics: HLT-NAACL 2004*, pages 145–152, Boston, Massachusetts, USA. Association for Computational Linguistics.
- Minh Phuong Nguyen and The Anh Le. 2024. [Feature-based unsupervised method for salient sentence ranking in text summarization task](#). In *Proceedings of the 2024 9th International Conference on Intelligent Information Technology, ICIIT '24*, page 346–351, New York, NY, USA. Association for Computing Machinery.
- Mihir Parmar, Hanieh Deilamsalehy, Franck Dernoncourt, Seunghyun Yoon, Ryan A. Rossi, and Trung Bui. 2024. [Towards enhancing coherence in extractive summarization: Dataset and experiments with LLMs](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 19810–19820, Miami, Florida, USA. Association for Computational Linguistics.
- Horacio Saggion and Thierry Poibeau. 2013. [Automatic text summarization: Past, present and future](#). In Thierry Poibeau, Horacio Saggion, Jakub Piskorski, and Roman Yangarber, editors, *Multi-source, Multilingual Information Extraction and Summarization*, pages 3–21. Springer Berlin Heidelberg, Berlin, Heidelberg.
- Tom Trabasso and Linda L. Sperry. 1985. [Causal relatedness and importance of story events](#). *Journal of Memory and Language*, 24(5):595–611.
- Salvatore Trani, Claudio Lucchese, R. Perego, David E. Losada, Diego Ceccarelli, and Salvatore Orlando. 2018. [SEL: A unified algorithm for salient entity linking](#). *Computational Intelligence*, 34:2–29.
- Florian Wolf and Edward Gibson. 2004. [Paragraph-, word- and coherence-based approaches to sentence ranking: A comparison of algorithm and human performance](#). In *Text Summarization Branches Out*, page 18, Barcelona, Spain. Association for Computational Linguistics.
- Amir Zeldes. 2017. [The GUM corpus: Creating multilayer resources in the classroom](#). *Language Resources and Evaluation*, 51(3):581–612.
- Amir Zeldes, Tatsuya Aoyama, Yang Janet Liu, Siyao Peng, Debopam Das, and Luke Gessler. 2025. [eRST: A signaled graph theory of discourse relations and organization](#). *Computational Linguistics*, 51(1):23–72.
- Amir Zeldes and Jessica Lin. 2026. What makes an entity salient in discourse? *Corpus Linguistics and Linguistic Theory*.

## A Annotation tool format

Our annotation tool, GlowLyter, is built for aligning text segments to content displayed in a carousel at the top of the interface, in this case containing five summaries per document. However the tool supports a generic format, shown in Figure 3, which incorporates a header (lines prefixed by #) declaring the desired annotation types.

In particular, the annos declaration is responsible for generating the annotation dialog elements seen in Figure 1. Check boxes are added using the syntax `annname:check`, numerical fields can be added by specifying a range of numbers (here: `duplicate:0-128`, since this text has 128 units), drop down lists can be added by specifying their values (`annname:val1,val2,val3,...`, not used above), and fields with no values are declared as simple key names (in this example, notes). The interface can therefore be used for a variety of annotation tasks.

The spans to be highlighted are given line by line below the header rows. For discontinuous discourse units we use the escape notation `\s-N`, where N is the content line that starts the discontinuous span. Thus `\s-2` on the fourth content line in the Figure indicates that this unit (unit 4) is part of the same unit as unit 2, i.e. the discontinuous span “*The lack of industry, .. compared to the North*,”. Other special codes can be used, such as `\b` to indicate a unit to be displayed in bold (for headings), `\i` for italics, `\l-NAME` for speaker labels and `\n` to add a line of space at the end of paragraphs.

## B Computing RST centrality

RST trees recursively divide a document into satellite and nucleus units connected by discourse relations. Satellites are considered less prominent than nuclei, and analysts assign satellite status to units that are less central to the pragmatic purpose or main meaning of a document or section. If two units form a coordination of equally important propositions, a multinuclear relation is postulated. For example, in the CONCESSION relation in Figure 4, the less important conceded unit [3-5] is a satellite of the nucleus, unit [2]. However Units [4] and [5] are coordinate and equally important.

```

# summary1 = This excerpt from a history book tells about the American Civil War, describing economic...
# summary2 = (human2) This section discusses some factors which led the Union to defeat the ...
# annos = approximate:check;component:check;duplicate:0-128;notes
15.4 The Union Triumphant
The lack of industry,
compared to the North,
\s-2 undercut the ability
to sustain and wage war.
...

```

Figure 3: Input format for the annotation interface

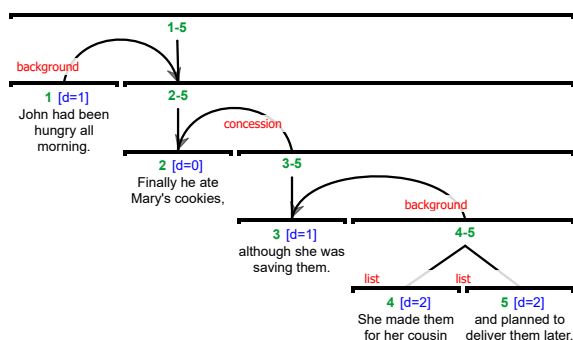


Figure 4: RST centrality example. Graph distance from the root is indicated in blue, e.g. [d=1] means one horizontal edge away from the root, which is unit [2].

The relative centrality of a unit can be measured based on the number of satellite relations (horizontal arrows) that are traversed between it and the document root, that is the unit (or units) which head a constituent that is not a satellite of any other constituent.

In the Figure, unit [2] is the root, which can be understood as the most salient unit from the RST perspective – John’s eating of the cookies is the most central event, and therefore has a distance of 0 ([d=0]). The constituents headed by [1] and [3] are direct satellites of the root, and therefore have [d=1]. Finally the coordination of [4] and [5] is a satellite of [3], and therefore has a greater distance, scoring [d=2].

### C RST vs. salience example

Figure 5 shows a fragment of an RST tree with nodes shaded by their aligned proposition salience. The most central discourse unit [3], which all units point to directly or indirectly, receives the maximum score of 5 (shaded red) because all summaries mention this proposition (*‘A sensitive .. document was found .. on an Ottawa street’*). This unit is considered a duplicate of the title (unit [1]), which is as salient, and likewise unit [5] receives the same score. The date of the event (*‘August 15, 2008’*)

is mentioned in only 3 summaries, corresponding to the orange highlight for unit [2]. Since no summary mentions unit [4] (the document being given to CBC), it receives no background color and remains white. Unit [6] is interrupted by unit [7] (not mentioned in any summary), and continued by unit [8], marked using the label SAME-UNIT.

Graph topology (graph distance from the most central or root unit, [3]) is correlated with salience, but not perfectly: for example the coordinate units [10] and [11] differ: [11] is mentioned by one summary (marked in yellow), but [10] is not.

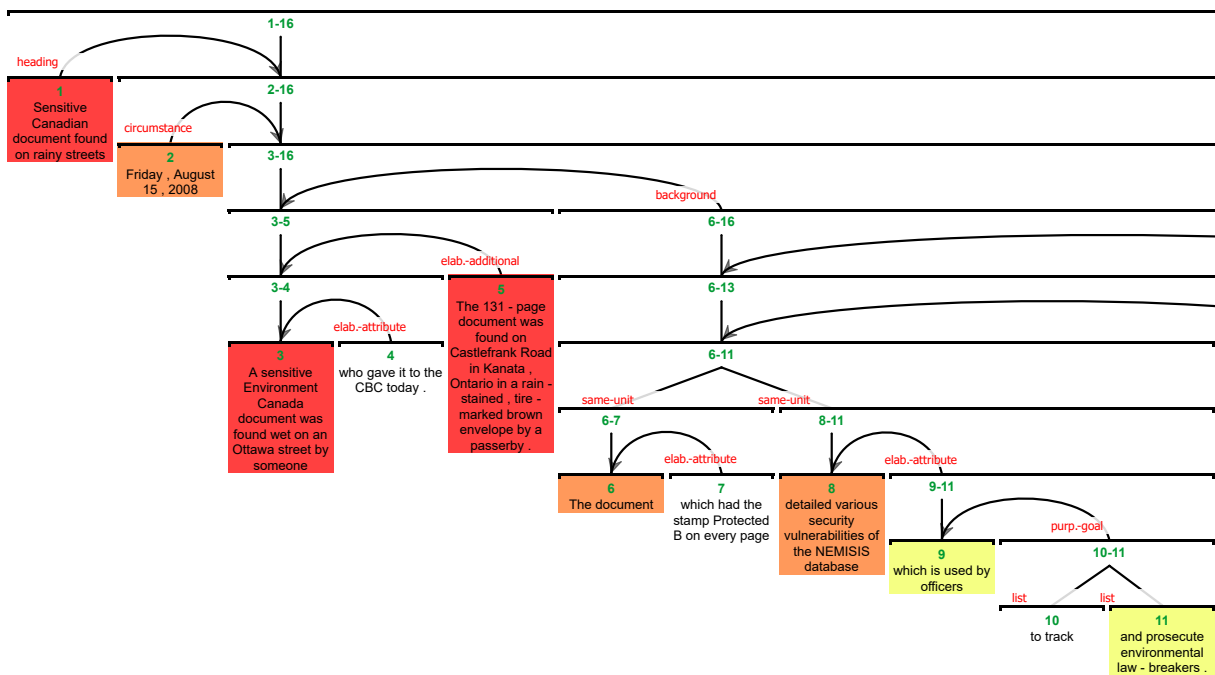


Figure 5: Fragment of an RST tree with units shaded by salience: a score of 5=red, 3=orange, 1=yellow.

# LLMs as annotators of credibility assessment in Danish asylum decisions: evaluating classification performance and errors beyond aggregated metrics

Galadrielle Humblot-Renaux<sup>1,2</sup>, Mohammad N. S. Jahromi<sup>1,3,2</sup>, Rohat Bakuri-Jørgensen<sup>1</sup>, Marieke Anne Heyl<sup>3</sup>, Asta S. Stage Jarlner<sup>3</sup>, Maria Vlachou<sup>4</sup>, Anna Murphy Høgenhaug<sup>3</sup>, Desmond Elliott<sup>4,2</sup>, Thomas Gammeltoft-Hansen<sup>3</sup>, Thomas B. Moeslund<sup>1,2</sup>

<sup>1</sup>Visual Analysis and Perception Lab, Aalborg University <sup>2</sup>Pioneer Center for AI, Denmark

<sup>3</sup>Center of Excellence for Global Mobility Law, University of Copenhagen

<sup>4</sup>Department of Computer Science, University of Copenhagen

Correspondence: [gegeh@create.aau.dk](mailto:gegeh@create.aau.dk)

## Abstract

Off-the-shelf large language models (LLMs) are increasingly used to automate text annotation, yet their effectiveness remains under-explored for underrepresented languages and specialized domains where the class definition requires subtle expert understanding. We investigate LLM-based annotation for a novel legal NLP task: identifying the presence and sentiment of credibility assessments in asylum decision texts. We introduce RAB-Cred, a Danish text classification dataset featuring high-quality, expert annotations and valuable metadata such as annotator confidence and asylum case outcome. We benchmark 21 open-weight models and 30 system-user prompt combinations for this task, and systematically evaluate the effect of model and prompt choice for zero-shot and few-shot classification. We zoom in on the errors made by top-performing models and prompts, investigating error consistency across LLMs, inter-class confusion, correlation with human confidence and sample-wise difficulty and severity of LLM mistakes. Our results confirm the potential of LLMs for cost-effective labeling of asylum decisions, but highlight the imperfect and inconsistent nature of LLM annotators, and the need to look beyond the predictions of a single, arbitrarily chosen model. The RAB-Cred dataset and code are available at <https://github.com/glhr/RAB-Cred>

## 1 Introduction

Deepening our understanding of the asylum-decision making process (e.g. discovering bias) requires understanding whether and how the applicant’s *credibility* is assessed across a large body of legal decisions. Credibility is a central element in almost all legal proceedings but is known to play an inordinate role in asylum cases, where adjudicators often find evidence to be scarce.

We specifically focus on the Danish asylum decision-making system. When the Danish Immigration Service rejects an asylum application,

the case is automatically appealed to the Danish Refugee Appeals Board (RAB) for reassessment. Written decision summaries spanning the past two decades, including the RAB’s legal reasoning, are publicly available and form the basis of this work.

When assessing an applicant’s eligibility for protection, the RAB often conducts a credibility assessment of the applicant’s testimony, that is, whether the applicant’s narrative of past events is deemed trustworthy and plausible. To date, credibility assessment in the Danish RAB’s decisions has only been studied at a small scale, relying on expert manual annotation (Rask Nielsen and Holten Møller, 2022; Høgenhaug et al., 2023; Hertz and Jarlner, 2025), a time-intensive process. Large Language Models (LLMs), as zero-shot or few-shot text classifiers, offer a potential path to labeling credibility assessments automatically, and thus enabling future large-scale studies in the legal domain.

However, identifying *whether* a credibility assessment was made and whether it was *positive* or *negative* is not straightforward. First, the process of credibility assessment is poorly understood and characterized by fuzziness, with no clear consensus among practitioners and researchers (Bendixen, 2020; Jarlner et al., 2026). Thus, the pre-existing general knowledge of off-the-shelf LLMs’ must be complemented with a detailed task definition provided by domain experts. Second, even with a specific definition, reliable annotation requires fine-grained semantic and contextual understanding of Danish legal texts.

In practice, due to the specificities of this data, difficult or atypical cases face a real risk of systematic misclassification by an LLM. Unlike the case outcome or demographic factors which can easily be extracted, annotating credibility assessment might require “reading between the lines”, as it is not necessarily stated explicitly (see Table 1). Asylum claims can be rejected despite a positive credibility assessment, or vice-versa, due to the im-

The Board granted a residence permit (Refugee Protection Status) to a female citizen of Somalia, born in 1989. She entered Denmark in February 2004. Like the Immigration Service, **the Board considered the applicant to be a Somali citizen. Based on the information in the case, the Board had to conclude that the applicant had never been to Somalia, did not speak the language, and had no family or other network in Somalia.** As the applicant was a single girl aged 16, the Board found that, based on the background information, it was probable that she would risk inhuman or degrading treatment covered by section 7(2) of the Aliens Act if she were deported to Somalia. The applicant had lived in Yemen since she was five years old without having had any conflicts with the Yemeni authorities. However, it was unclear on what basis she had been residing in Yemen. As it was therefore uncertain whether the applicant had legal residence in Yemen and could enter Yemen, the majority of the Board found that Yemen could not be considered the applicant’s first country of asylum pursuant to section 7(3) of the Aliens Act. As a result, the Refugee Appeals Board granted the applicant a residence permit pursuant to section 7(2) of the Aliens Act.

(Original case text: <https://fln.dk/praksis/2019/april/somalia-somalia20054/>)

Q1: Credibility assessment present?	Yes (Confidence: Medium)
Q2: Credibility assessment sentiment	Positive (Confidence: High)

Table 1: Translation of the shortest written decision from the validation set, and its corresponding annotation agreed upon by the two domain experts. Translated with DeepL.com (free version). For reference, we highlight the extract most indicative of a positive credibility assessment for domain experts.

portant distinction between credibility assessment (are facts accepted?) and risk assessment (are there sufficient grounds for asylum, given background material and accepted facts?). These two assessments can be difficult to disentangle, as they are often proximal in articulation and position within the same text; small differences in phrasing can change the final label entirely. Moreover, a single decision can also contain elements pointing to both positive and negative credibility, when some facts are accepted while others are rejected. These challenges are compounded by the linguistic setting: RAB decisions are written in Danish, a medium-resource language, using specialized legal terminology.

Domain experts have an interest not only in LLM annotations being as accurate as possible, but also in understanding what *types* of error occur, for *which cases*, and whether these mistakes are *understandable* or *unacceptable*. Our aim is therefore to investigate the extent to which the annotation of credibility assessment in Danish asylum decisions can be reliably automated by off-the-shelf LLMs, with a particular focus on annotation error. Our contributions are summarized as follows:

- we present RAB-Cred, an expertly annotated Danish legal text classification dataset from an under-represented domain and language, which poses interesting challenges for legal experts and natural language understanding.
- to explore the potential of zero-shot and few-shot classification for this task, we systematically benchmark 21 open-weight multilingual LLMs across 30 different prompts.
- beyond standard aggregated performance met-

rics, we analyse the errors produced by and across top-performing models, relating them to human label confidences.

The RAB-Cred dataset (including expert multi-annotator labels, and self-reported confidence, and case outcome), along with LLM annotations and code to reproduce the analysis are available at <https://github.com/glhr/RAB-Cred>.

## 2 Related work

**Task and dataset** There is a growing interest in treating LLMs as a drop-in replacement for human annotators/coders in social sciences and humanities (Törnberg, 2024; Ziems et al., 2024; Davidson, 2024; Halterman and Keith, 2025; Wen et al., 2025; Meizlish and Ziffo, 2026). In legal research specifically, off-the-shelf LLMs have been evaluated on a variety of tasks ranging from argument mining (Held and Habernal, 2025) to legal interpretation classification (Dugac and Altwicker, 2025), showing mixed results. To the best of our knowledge, the identification and sentiment classification of credibility assessment in legal decisions constitutes a novel NLP task, absent from existing legal text classification datasets (Ariai et al., 2025).

Furthermore, publicly available and annotated texts datasets within refugee law are especially scarce - AsyLex (Barale et al., 2023) currently being the only one to our knowledge. Unlike AsyLex, RAB-Cred contains non-English texts. Compared to English, Danish is under-represented in LLM training corpora (Enevoldsen et al., 2023; Zhang et al., 2024; Ekgren et al., 2024) and evaluation benchmarks (Singh et al., 2025; Ahuja et al., 2023; Xuan et al., 2025).

**Models and prompts** We include a wide open-weight model selection, including recent multilingual releases such as EuroLLM 22B and Bielik 11B v3 (Ramos et al., 2026; Ociepa et al., 2025b) - this contrasts to the majority of related studies which only consider a small handful of (often proprietary) models (Pavlovic and Poesio, 2024). The importance of a well-crafted prompt is repeatedly highlighted in the LLM annotation literature, but the effectiveness of different formulations is largely model, domain and task-specific (Pavlovic and Poesio, 2024; Mizrahi et al., 2024). We therefore treat prompt choice as a key, model-specific hyperparameter in our experiments. We benchmark state-of-the-art approaches including chain-of-thought (CoT), metacognitive and few-shot prompting (Vatsal and Dubey, 2024), and also experiment with providing varying levels of detail and context, similarly to Majer and Šnajder (2024).

**Mistakes matter** In order to reflect on the limitations of LLM annotators and on the data itself, we go beyond aggregated classification performance of individual models and zoom into how and when mistakes occur. In a similar spirit, Halterman and Keith (2025) take a holistic approach when evaluating LLM’s ability to follow a codebook for annotation, including manual error analysis and identification of shortcut behaviour. We go a step further and also look at consistency and correctness of LLM-generated annotations *across* models and prompts, at the level of individual samples. Our evaluation draws from error and prompt sensitivity analyses performed in existing work (Majer and Šnajder, 2024; Zhuo et al., 2024), but we instead consider an ensemble of 15 top-performing LLM annotators. Moreover, input from domain experts enables us to relate LLM misclassifications to human confidence, and to assess their severity.

### 3 The RAB-Cred dataset

The dataset used in this work is based on public asylum decisions made by the Danish Refugee Appeals Board (RAB), available at <https://fln.dk/praksis>. The RAB is the second and final legal instance of the Danish asylum system. Cases rejected by the Immigration Service are automatically appealed to the Board. Written decisions are relatively brief (around 600 words on average) - although some extend beyond 1500 words, cf. Figure 1. They first outline details about the case and asylum motives, and then explain the RAB’s deci-

sion to either uphold or over-turn the Immigration Service’s rejection, or to remand the case. Additional information about the data, metadata and annotations can be found in Appendix A.

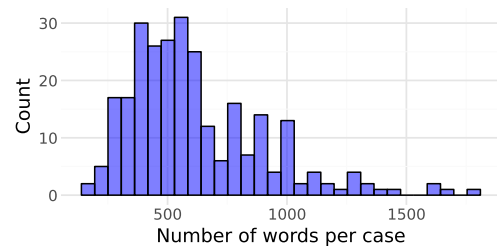


Figure 1: Distribution of case lengths in RAB-Cred.

Given a written decision, the aim is to identify *whether a credibility assessment was made*, and if so, whether credibility was assessed *positively or negatively*.

#### 3.1 Validation and test sets

We sample 273 cases from the RAB’s official website using stratified random sampling across time ranges of interest (cf. Appendix A for details). 73 of these cases are used as a *validation* set, for the human annotators to iteratively develop a codebook and for selecting optimal model-prompt combinations when generating LLM annotations. The remaining 200 cases are held out as an *unseen test* set, which we use to quantify human inter-annotator agreement and analyse the errors made by LLMs.

#### 3.2 Codebook and class definition

Unlike Drápal et al. (2023); Bay-Jørgensen et al. (2026) who involve LLMs in designing the annotation task itself (e.g. identifying relevant concepts to annotate), we rely on a codebook developed by refugee law experts, which defines the categories and annotation guidelines. This codebook is used as a basis for both human annotation and LLM prompts - an approach which has shown promising results in recent work (Ruckdeschel, 2025; Halterman and Keith, 2025).

Specifically, two domain experts and annotators (nicknamed H1 and H2) set out to jointly define what exactly constitutes a credibility assessment in the RAB dataset, and how it should be annotated at the text level. Through discussions and interactive annotation sessions, the codebook was iteratively refined until full inter-annotator agreement was reached on the validation set. The annotators converged to a 3-tiered categorization, where the credibility assessment is annotated as either Absent,

Positive, or Negative. For each case, the annotators also recorded their confidence level (Low, Medium or High) about the presence of a credibility assessment, and its sentiment (if present). Note that these are posed as two separate questions, as an annotator may be highly uncertain about whether certain statements qualify as a credibility assessment, but may be highly confident that *if* they qualify as a credibility assessment, the assessment is positive. Table 1 shows an example annotation.

Figure 2 shows that according to the two experts, a credibility assessment is present in over 75% of cases; when present, it is more likely to be negative. Looking at annotator confidence suggests that while the experts are largely confident in their annotation, the *presence* of a credibility assessment is more difficult to ascertain than its sentiment.

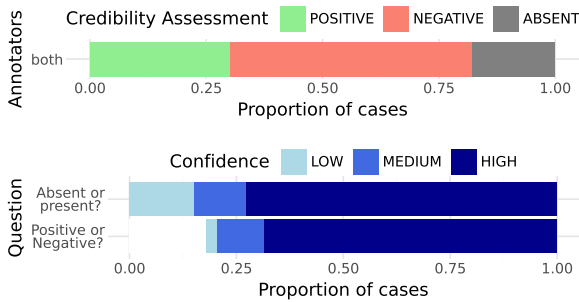


Figure 2: Label (top) and confidence (bottom) distribution of human annotators on the validation set.

### 3.3 Inter-annotator agreement on the test set

Gold-standard annotation of the test set was performed by the same two domain experts H1 and H2 as for the validation set, but completely independently. We observed a very high level of agreement on the presence and sentiment of a credibility assessment (Cohen’s  $\kappa = 0.97$ ), with only 4 cases (out of 200) for which the annotators diverge - as shown in Figure 3. Interestingly, for 3 out of these 4 cases, both annotators reported high confidence. The high inter-rater agreement suggests that the annotations are of high quality and that the codebook is sufficiently informative, but the presence of high-confidence disagreement suggests that even with a strong, detailed understanding of the codebook, some ambiguity or possible conflicting interpretation remains.

As shown in Figure 4, when independently annotating test samples, both annotators are more likely to be unsure about the presence of a credibility assessment than its sentiment.

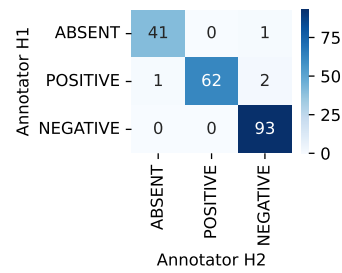


Figure 3: Confusion matrix showing inter-annotator agreement on the test set (H1 vs. H2’s annotations).

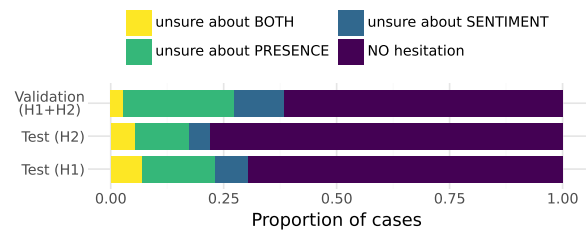


Figure 4: Inter-class hesitation of the human annotators, inferred from annotator confidence.

To resolve disagreement between H1 and H2, a third domain expert H3 (who was not involved in codebook development and was not shown H1 and H2’s labels) annotated the 4 cases; for all 4, H3 aligned with either H1 or H2. The final gold label is taken as the majority vote.

### 3.4 Correlation with outcome

Figure 5 shows a notable relation between case outcome and credibility: 80.6% of reversed cases in RAB-Cred (asylum granted) are associated with a positive credibility assessment, and 66.3% of upheld rejections with a negative one. At the same time, for 41.5% of cases, either no credibility assessment is present or the outcome contradicts the credibility sentiment. The outcome should therefore not be used as a proxy for annotation.

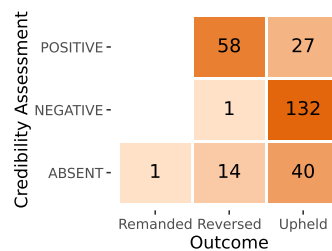


Figure 5: Relation between case outcome and gold credibility assessment labels on the full dataset.

## 4 LLM-generated annotations

Given the small amount of human-labeled data at our disposal, our annotation pipeline relies on off-the-shelf LLMs used as zero-shot or few-shot classifiers. We first describe our selection of language models, along with the prompt variants used to evaluate automatic credibility annotations. We then systematically evaluate the role of model choice and prompt choice for this task. The full list of models and implementation details are described in Appendix B.2, the prompts and the few-shot examples are in Appendix B.1, and detailed results in Appendix C.

### 4.1 Models and configuration

When selecting potential off-the-shelf LLMs, we considered a wide range of instruction-tuned models and applied the following criteria:

1. **Open-weight** to ensure reproducibility and also to enable offline annotation, which is necessary for annotating sensitive data.
2. Explicitly trained on **multilingual** data.
3. **Limited model size** due to compute constraints (a single H100 with 80 GB VRAM).
4. **Context length** of at least 8K tokens, to accommodate long case texts, detailed instructions and LLM reasoning output.

This resulted in 21 models spanning 9 model families (Gemma (Team et al., 2025), Qwen (Team, 2024), Phi (Abdin et al., 2024), Aya (Dang et al., 2024), Granite (IBM, 2025), Mistral (Liu et al., 2026), Llama (Grattafiori et al., 2024), Bielek (Ociepa et al., 2025a) and EuroLLM (Ramos et al., 2026)), ranging from 3B to 35B parameters.

### 4.2 Prompt variants

We separately consider the role of the *system* prompt and the *user* prompt for annotating credibility. The system prompt is used to provide background knowledge (e.g. what is a credibility assessment?), while the user prompt determines how the classification task should be tackled (e.g. directly giving the final answer, using examples or following multiple steps).

We design 6 system prompts and 5 user prompts, yielding a total of 30 unique prompts templates. All prompts are written in English with Danish case text as input, as Lai et al. (2023); Pavlovic and Poesio (2024) found this approach effective for non-English data.

**System prompts (SP)** System prompts (SP) follow a nested hierarchy of increasing domain context. SP0 provides no system prompt (baseline). SP1 assigns a domain-expert persona. SP2 extends SP1 with the verbatim codebook. SP3 restructures SP2 with indicative Danish phrases per class. SP4 extends SP3 with critical edge cases (e.g., hypothetical legal constructions, mixed sentiment). SP5 extends SP1 by offering an alternative expert-written breakdown that explicitly disambiguates credibility from risk assessment and neutral reporting.

**User prompts (UP)** We design 5 user prompts of increasing complexity. UP1 directly instructs the model to select one of three classes. UP1-FS extends UP1 with three labeled examples (selection described below). UP2 decomposes the task into two binary questions: first whether a credibility assessment is present, then its sentiment, mirroring the human annotation structure. UP3 and UP4 introduce an unconstrained reasoning step before the final classification: UP3 via zero-shot CoT prompting (Kojima et al., 2022) and UP4 via zero-shot metacognitive prompting (Wang and Zhao, 2024).

**Few-shot examples** We select one example per class from the validation set (thus three examples in total), choosing cases that are unambiguous for domain experts (high label confidence) yet consistently challenging for LLMs (highest zero-shot misclassification rate in preliminary experiments).

From a domain perspective, each example represents an atypical scenario: the "absent" case lacks a credibility assessment; the "positive" and "negative" cases each exhibit a mismatch between credibility sentiment and outcome.

### 4.3 Comparing models and prompts

We first investigate which models and prompting strategies are best suited for annotating credibility assessment. We evaluate the 21 selected models  $\times$  30 user-system prompt combinations on the validation set - excluding the 3 few-shot samples used in UP1-FS, leaving 70 samples.

Given class imbalance, we report macro F1 score as the primary metric for classification performance. As a lower baseline, the outcome is used as a naive heuristic for credibility assessment using the following mapping: remanded  $\rightarrow$  absent, reversed  $\rightarrow$  positive, upheld  $\rightarrow$  negative.

**Comparing models** Figure 7 shows the size-performance tradeoff considering only the best

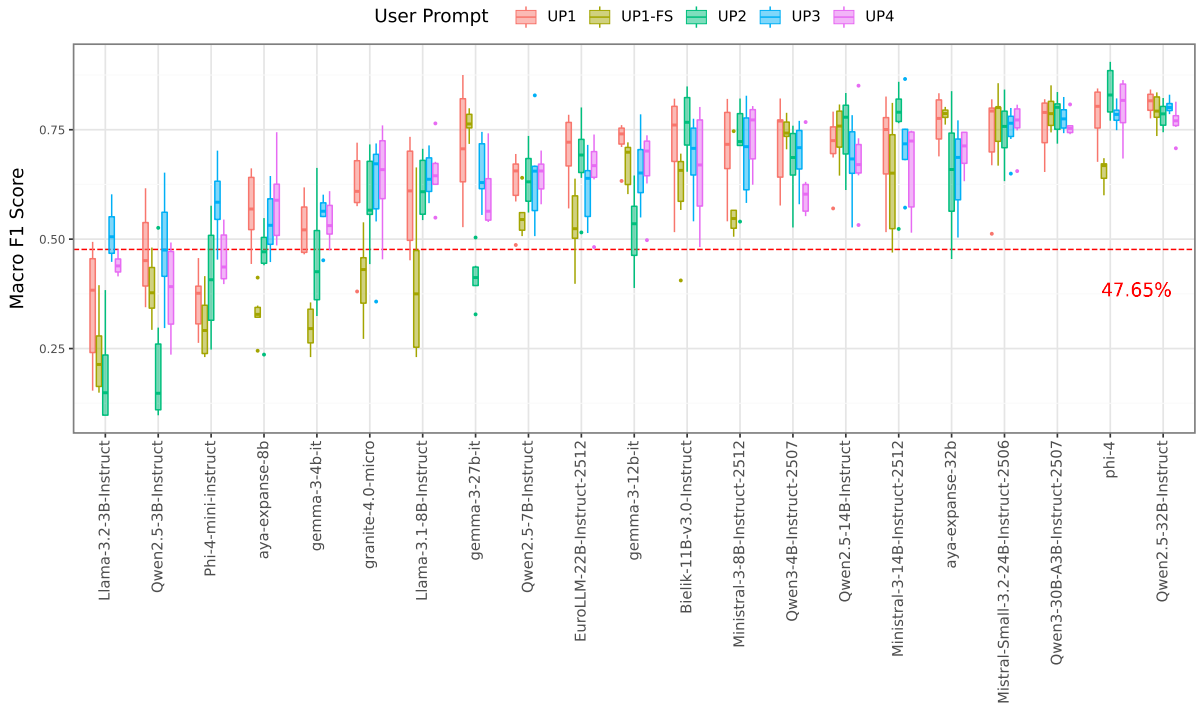


Figure 6: Validation set classification performance per model for different user prompts. Each boxplot is across 6 system prompts. Models are ordered by average macro-F1 score. Red line: outcome-as-credibility baseline.

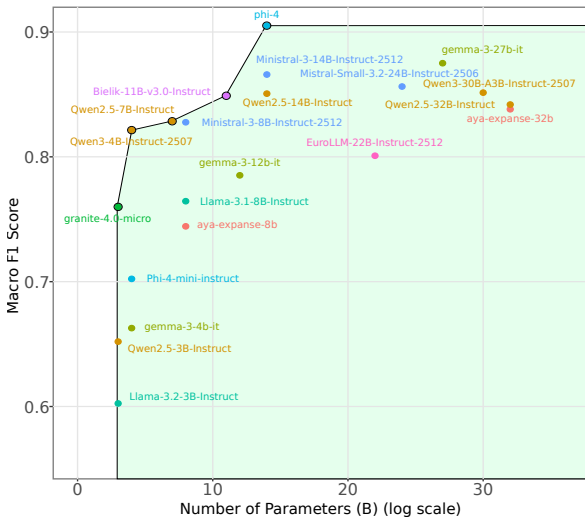


Figure 7: Best-case classification performance (taking the top 1 system-user prompt for each model) on the validation set, as a function of model size.

system-user prompt per model, while Figure 6 shows model-specific performance across all prompt variants. With the exception of Qwen2.5-7B and Ministral-3-8B, all models that consistently outperform the outcome-as-credibility baseline are larger than 11B. However, as shown in Figure 7, increases in size do not necessarily translate to performance gains. Among models under 10B, Qwen2.5-

7B and Qwen3-4B offer the best trade-off between size and performance. Perhaps most strikingly, phi-4 (14B) achieves the highest F1 on the validation set (90.51%), while the largest models in our selection, including Qwen2.5-32B, Qwen3-30B and aya-expense-32b, plateau at or below 85% F1.

As shown in Figure 8, models vary considerably in how sensitive they are to different prompt combinations. Models such as Qwen2.5-32B are highly robust, while performance for a model like Gemma-3-27B can vary by over 54% in F1. A general tendency is that bigger models vary less in performance with prompt changes. In the absence of any explanation of what a credibility assessment is (SP0 or SP1), the 2 largest Qwen models and phi-4 stand out as the best performing models, with Qwen2.5-32B exceeding 83% F1 using UP1 both with (SP1) and without (SP0) a persona.

**Comparing system prompts** Looking at Figure 8, although the relative effectiveness of different SPs is UP-dependent, there is a clear benefit in incorporating expert knowledge (SPs 2-5), with SP3 and SP4 showing the most consistent performance. This shows the benefit of interdisciplinary prompt design: different from SP2 and SP5 which were written by domain experts alone, SP3 and SP4 were formulated by computer scientists based

on the codebook and with domain expert feedback. Adding SP4’s edge-cases to SP3’s handcrafted context is effective in some cases, but not on average.

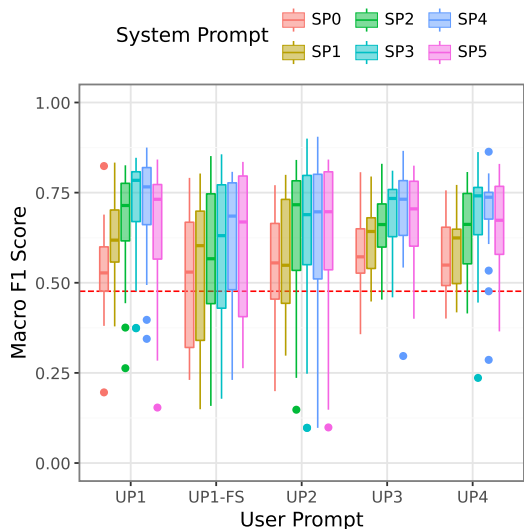


Figure 8: Validation set classification performance across 21 models for different user-system prompt combinations. Red line: outcome-as-credibility baseline.

**Comparing user prompts** Compared to the basic UP1 prompt, the effect of more advanced prompting strategies (few-shot, multi-turn, and CoT/metacognitive) is found to be highly model-specific. For instance, while Phi-4 greatly benefits from having the classification task broken down into 2 binary questions with UP2 (such that it achieves the highest performance of any model-prompt combination), this approach heavily degrades performance for Gemma models. We find few-shot prompting (UP1-FS) to be effective only for Qwen3-30B, Mistral-Small and Qwen2.5-4B. As for reasoning-based prompts (UP3 and UP4), these reduce the variation in performance across SPs and models (Figure 8). CoT prompting only outperforms UP1 when the SP lacks sufficient context (SP0 and SP1.)

**Summary** Overall, while adding explanations of credibility assessment through the system prompt clearly helps, no single prompt emerges as the “winner”, which confirms the need for model-specific prompt choice. In the zero-shot setting, Phi-4 shows impressive performance for its size, see Figure 7. This is somewhat surprising, given that it is seldom evaluated in the LLM-as-annotator literature, and its model card states “multilingual data constitutes about 8% of our overall data” and “phi-4 is not intended to support multilingual use”.

## 5 When and how do the best models fail?

Our aim is to analyze classification errors and agreement among top-performing LLM annotators. We select the most promising model-prompt combinations based on performance on the validation set. Since the optimal prompt is highly model-dependent, we rank each model according to its average macro-F1 score across its top-3 prompts, then select the top 5 models (phi-4, gemma-3-27b-it, Ministral-3-14B-Instruct-2512, Ministral-3-14B-Instruct-2512, Qwen3-30B-A3B-Instruct-2507 - each paired with its 3 highest-performing prompts) for final evaluation on the test set. The resulting selection contains 11 unique user-prompt combinations, with SP5+UP2 being the most frequent (3 instances). Individual classification performance for these 15 model-prompt combinations is reported in Appendix C and D.0.1.

When compared to the majority label assigned by human annotators, the macro F1-Score of the selected model-prompt combinations ranges from 84.4% to 94.7% on the test set, with phi-4 being the strongest model, and gemma-3-27b-it being the weakest on average. For reference, the outcome-as-classifier baseline achieves an F1-Score of 53%.

	annotator pair	agreement with human majority	inter-annotator agreement
Domain experts	H1	0.984	0.967
	H2	0.983	
Mistral-Small-3.2-24B-Instruct-2506	SP3+UP1-FS	0.883	0.922
	SP4+UP1	0.882	
phi-4	SP4+UP2	0.906	0.913
	SP4+UP4	0.913	

Table 2: Cohen’s  $\kappa$  for 3 annotator pairs on the test set.

**Inter-LLM agreement** We compare the human inter-annotator agreement on the test set (between H1 and H2) with that of LLM annotators in Table 2. We select two pairs of LLM annotators with near-identical agreement levels with respect to the human majority label (i.e. similar level of “correctness”), and measure the agreement within the pair. We find the LLM annotator pairs to be less aligned than the human annotator pair, even when they share the same model weights and system prompt, indicating that misclassifications are not necessarily consistent across similarly-performing models. The inter-annotator agreement for all possible LLM annotator pairs can be found in Appendix D.0.2. The highest level of agreement ( $\kappa = 0.950$ ) occurs between two Gemma models with the same user prompt, but different system prompts (SP3 vs. SP4), despite their F1 differing by almost 3%.

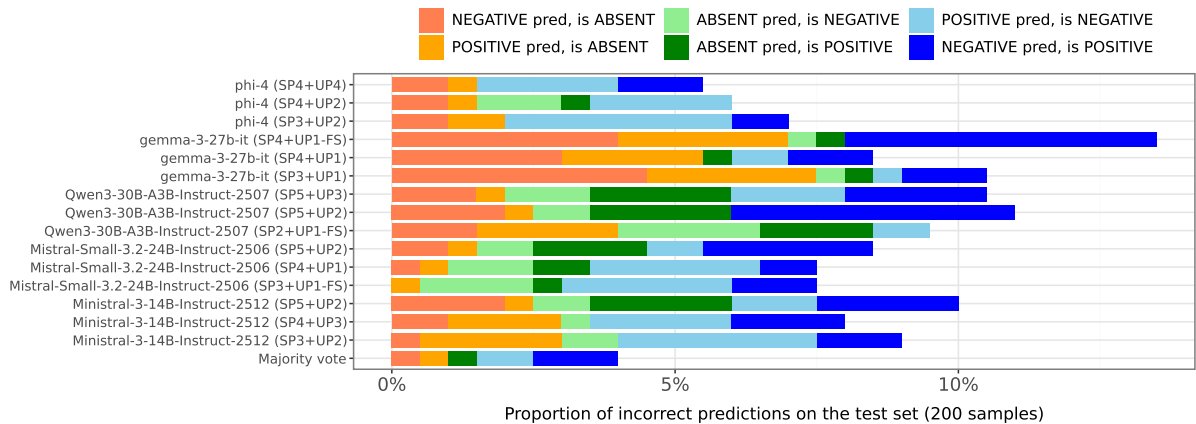


Figure 9: Individual LLM and ensemble mistakes, color-coded by class confusion (cf. Appendix D.0.4).

**Instance-level sensitivity** Prompt sensitivity is often measured at the dataset level, looking at how aggregate performance measures vary across different prompts (Hua et al., 2025; Mizrahi et al., 2024), similarly to Section 4.3. However, two LLM annotators with the same performance may be misclassifying different instances. To complement our aggregate analysis, we adopt the PromptSensiScore (PSS) from Zhuo et al. (2024), which instead operates at the instance level: it captures changes in the correctness of individual predictions, given a change of prompt. To compare the effect of model vs. prompt choice, we apply PSS in two ways: fixing the model and varying the prompt vs. fixing the prompt and varying the model (detailed results in Appendix D.0.3).

Examining prompt sensitivity first, we find that phi-4, gemma-3-27b, and Mistral-Small-24B all exhibit lower instance-level sensitivity than Ministral and Qwen3. Interestingly, Qwen3-30B shows the highest sensitivity despite strong aggregate stability (F1 score variation across prompts), suggesting its consistent performance masks instability on specific instances. This pattern underscores that dataset-level and instance-level metrics capture complementary axes of robustness.

For model sensitivity, we select a fixed prompt (SP5+UP2) shared across Ministral-14B, Mistral-Small-24B, and Qwen3-30B, and look at instability related to model change. The resulting PSS score of 0.05 is lower than model-specific PSS scores (Qwen: PSS $\approx$ 0.10, Ministral: PSS $\approx$ 0.08, Mistral: PSS $\approx$  0.06), indicating that prompt variations introduce greater instability than model differences, and underscoring the importance of prompt design.

**Inter-class confusion** In Figure 9, we zoom into the number and types of mistakes made by the 15 individual LLM annotators and by the ensemble (taking the majority vote). For individual LLMs, we find that the relative prevalence of different types of mistakes varies significantly across both models and prompts. Some LLMs never miss the presence of a credibility assessment (green in Figure 9), however all LLMs falsely identify a credibility assessment at least once (orange/salmon). Distinguishing between positive vs. negative credibility assessments seems less straightforward for LLM annotators than for human annotators (cf. Figure 4): when taking the majority vote, over half of the mistakes are *sentiment* misclassifications.

**Fine-grained analysis** The small number of test set samples allows us to visualize the correctness and agreement of LLMs on a case-by-case basis, as shown in Figure 10. 72% of cases (144 out of 200), are correctly classified by all 15 LLM annotators, 95% (190 out of 200) are correctly classified by at least half (i.e. at least 8 LLM annotators out of 15). Interestingly all the cases correctly classified by less than 75% of LLM annotators were assigned medium or low confidence by one of the annotators.

When taking the majority vote across LLM annotators, 96% of cases are correctly classified - 1.5 percentage points above the top single-model accuracy. For the remaining 8 misclassified cases (cf. Appendix D.0.4), we asked domain expert H1 to judge the LLM majority prediction: H1 considers the misclassification "acceptable" in 4 cases, "understandable" in 2 cases and "unacceptable" in 2 cases. Further details about this categorization are in Appendix D.0.5. We find that *all* model-prompt combinations make at least 1 unacceptable mistake.

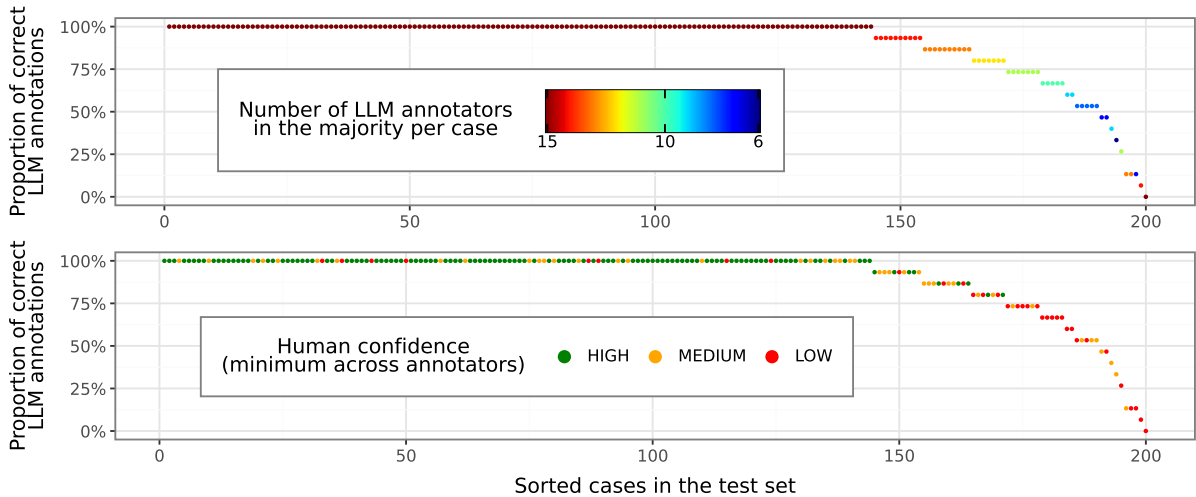


Figure 10: LLM agreement vs. LLM correctness vs. human confidence. Each point corresponds to a single case in the test set, and the points in both plots are sorted by correctness (identical ranking in both plots).

Lastly we zoom into the two cases which were misclassified by 14 or all 15 LLM annotators respectively (bottom right of the plots in Figure 10): both were labeled as having no credibility assessment by H1 and H2. Interestingly, H1 considers the LLM majority prediction to be acceptable in both cases, to the point of reconsidering their own annotation. In one case, hesitation is due to the use of language typically associated with credibility assessments but a future-oriented judgment, while in the second case, credibility assessment of the claimant’s relative may be misconstrued with the claimant’s own credibility. We refer to Appendix D.0.5 for the two case texts, LLM reasoning and domain expert reasoning.

## 6 Discussion and future work

From a practical standpoint, our results serve as a solid baseline for the RAB-Cred dataset and suggest that automating the annotation of credibility assessment using LLMs is a promising direction, but not a perfect replacement for manual expert annotation. The case-by-case analysis provides preliminary support for the use of prompt and model ensembling, as LLM aggregation yields a clear improvement over any single LLM annotator. Furthermore, ensembling of LLM annotators could enable a human-in-the-loop approach where cases with high inter-model or inter-prompt disagreement are flagged for expert review, while the rest are annotated automatically. A systematic comparison of ensembling strategies and their cost-performance trade-off is a natural direction for future work.

A consistent finding across our experiments is that prompt design matters at least as much as model choice. The effectiveness of advanced prompting strategies is highly model-dependent, and at the instance level, changing the prompt can affect individual predictions more than changing the model. This confirms the need for multi-prompt evaluation in LLM annotation studies (Mizrahi et al., 2024). It also suggests that multi-prompt ensembling with a single model could be sufficient.

Furthermore, several design choices in our study point to avenues for further investigation. We use English-language prompts on Danish input texts, following evidence that this cross-lingual approach is effective for multilingual classification (Lai et al., 2023; Pavlovic and Poesio, 2024). However, the interaction between prompt language, input language, and model choice is underexplored for legal texts specifically; prompting in Danish or translating case texts to English may yield different error profiles. Similarly, we evaluate only off-the-shelf, general-purpose LLMs. Domain-specialized or fine-tuned models, as advocated by Dominguez-Olmedo et al. (2025), may achieve higher accuracy, particularly on the boundary cases identified in our error analysis—though this comes at the cost of requiring labeled training data, which our approach aims to circumvent. While we focus on the final classification output, systematically analysing the intermediate reasoning traces produced by chain-of-thought and metacognitive prompts could offer further insight into *how* models arrive at their classification and whether this aligns with expert-provided explanations.

## Limitations

Our evaluation is based on a dataset of 273 cases, of which 3 are used as few-shot examples, 70 as a validation set, and 200 as a test set. While carefully annotated by experts, this dataset is not representative of the full body of publicly available Danish RAB decisions, which spans over 10,000 cases. The relatively small size limits the statistical power of comparisons between model-prompt combinations and may not capture the full diversity of credibility assessment formulations found in practice.

In terms of model selection, we restrict our evaluation to open-weight models of at most 35B parameters, due to compute constraints and the practical need for offline inference when working with sensitive legal data. We do not evaluate closed-source models such as GPT-4o or Claude, which may achieve stronger performance but cannot be deployed locally. Our findings therefore characterize the current capabilities of open-weight, moderately-sized LLMs, and should not be taken as an upper bound on what LLMs can achieve for this task.

Furthermore, all results are based on a single run per model-prompt combination. While we use greedy decoding throughout, which is deterministic for a given input, we do not account for possible stochasticity arising from numerical precision or hardware differences.

Whenever possible, we used constrained decoding via the *outlines* library to ensure that models produce valid category labels (cf. Appendix B.2.2 for details). Despite its practical advantages, constrained decoding is known to potentially degrade generation quality (Schall and de Melo, 2025), and its interaction with classification performance across different model architectures has not been systematically studied here.

Finally, although our stratified sampling partially addresses temporal variation in RAB decisions, we do not analyse whether model performance varies across time periods. Changes in legal practice, writing conventions, or anonymization procedures over the two decades covered by the dataset may introduce systematic differences that have not been captured in the evaluation.

## Acknowledgments

This work was supported by the Villum Foundation (“XAI-CRED”, grant no. 69198), the Grundfos Foundation (“REPAI”, grant no. 83648813), and the Danish National Research Foundation (“Center of Excellence for Global Mobility Law”, grant no. DNRF169).

Part of the computation done for this project was performed on the UCloud interactive HPC system, which is managed by the eScience Center at the University of Southern Denmark. Part of the computation was also performed on the AI Cloud HPC system managed by CLAUDIA at Aalborg University.

## References

- Marah Abdin, Jyoti Aneja, Harkirat Behl, Sébastien Bubeck, Ronen Eldan, Suriya Gunasekar, Michael Harrison, Russell J. Hewett, Mojan Javaheripi, Piero Kauffmann, James R. Lee, Yin Tat Lee, Yuanzhi Li, Weishung Liu, Caio C. T. Mendes, Anh Nguyen, Eric Price, Gustavo de Rosa, Olli Saarikivi, and 8 others. 2024. [Phi-4 technical report](#). *Preprint*, arXiv:2412.08905.
- Kabir Ahuja, Harshita Diddee, Rishav Hada, Millicent Ochieng, Krithika Ramesh, Prachi Jain, Akshay Nambi, Tanuja Ganu, Sameer Segal, Mohamed Ahmed, Kalika Bali, and Sunayana Sitaram. 2023. [MEGA: Multilingual evaluation of generative AI](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 4232–4267, Singapore. Association for Computational Linguistics.
- Farid Ariai, Joel Mackenzie, and Gianluca Demartini. 2025. [Natural language processing for the legal domain: A survey of tasks, datasets, models, and challenges](#). *ACM Comput. Surv.*, 58(6).
- Claire Barale, Mark Klaisoongnoen, Pasquale Minervini, Michael Rovatsos, and Nehal Bhuta. 2023. [AsyLex: A dataset for legal language processing of refugee claims](#). In *Proceedings of the Natural Legal Language Processing Workshop 2023*, pages 244–257, Singapore. Association for Computational Linguistics.
- Anna Bavaresco, Raffaella Bernardi, Leonardo Bertolazzi, Desmond Elliott, Raquel Fernández, Albert Gatt, Esam Ghaleb, Mario Giulianelli, Michael Hanna, Alexander Koller, Andre Martins, Philipp Mondorf, Vera Neplenbroek, Sandro Pezzelle, Barbara Plank, David Schlangen, Alessandro Suglia, Aditya K Surikuchi, Ece Takmaz, and Alberto Testoni. 2025. [LLMs instead of human judges? a large scale empirical study across 20 NLP evaluation tasks](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics*

- (Volume 2: Short Papers), pages 238–255, Vienna, Austria. Association for Computational Linguistics.
- Frederik Bay-Jørgensen, Henrik P Olsen, Mohammad N Jahromi, B. M Thomas, and Thomas Gammeltoft-Hansen. 2026. [Managing fuzziness: Leveraging llms for discovering credibility indicators in asylum cases](#). *Preprint*, SocArXiv:r29fv\_v1.
- Michala Clante Bendixen. 2020. [WELL-FOUNDED FEAR – CREDIBILITY AND RISK ASSESSMENT IN DANISH ASYLUM CASES](#). Refugees Welcome.
- John Dang, Shivalika Singh, Daniel D’souza, Arash Ahmadian, Alejandro Salamanca, Madeline Smith, Aidan Peppin, Sungjin Hong, Manoj Govindassamy, Terrence Zhao, Sandra Kublik, Meor Amer, Viraat Aryabumi, Jon Ander Campos, Yi-Chern Tan, Tom Kocmi, Florian Strub, Nathan Grinsztajn, Yannis Flet-Berliac, and 26 others. 2024. [Aya expanse: Combining research breakthroughs for a new multilingual frontier](#). *Preprint*, arXiv:2412.04261.
- Thomas Davidson. 2024. [Start generating: Harnessing generative artificial intelligence for sociological research](#). *Socius*, 10:23780231241259651.
- Ricardo Dominguez-Olmedo, Vedant Nanda, Rediet Abebe, Stefan Bechtold, Christoph Engel, Jens Frankenreiter, Krishna P. Gummadi, Moritz Hardt, and Michael Livermore. 2025. [Lawma: The power of specialization for legal annotation](#). In *The Thirteenth International Conference on Learning Representations*.
- Jakub Drápal, Hannes Westermann, and Jaromir Savelka. 2023. [Using large language models to support thematic analysis in empirical legal studies](#). *Preprint*, arXiv:2310.18729.
- Igor Dugac and Tilmann Altwicker. 2025. [Classifying legal interpretations using large language models](#). *Artificial Intelligence and Law*.
- Ariel Ekgren, Amaru Cuba Gyllensten, Felix Stollenwerk, Joey Öhman, Tim Isbister, Evangelia Gogoulou, Fredrik Carlsson, Judit Casademont, and Magnus Sahlgren. 2024. [GPT-SW3: An autoregressive language model for the Scandinavian languages](#). In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 7886–7900, Torino, Italia. ELRA and ICCL.
- Kenneth Enevoldsen, Lasse Hansen, Dan S. Nielsen, Rasmus A. F. Egebæk, Søren V. Holm, Martin C. Nielsen, Martin Bernstorff, Rasmus Larsen, Peter B. Jørgensen, Malte Højmark-Bertelsen, Peter B. Vahlstrup, Per Møldrup-Dalum, and Kristoffer Nielbo. 2023. [Danish foundation models](#). *Preprint*, arXiv:2311.07264.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, Amy Yang, Angela Fan, Anirudh Goyal, Anthony Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev, Arthur Hinsvark, and 542 others. 2024. [The llama 3 herd of models](#). *Preprint*, arXiv:2407.21783.
- Andrew Halterman and Katherine A. Keith. 2025. [Codebook llms: Evaluating llms as measurement tools for political science concepts](#). *Political Analysis*, page 1–17.
- Lena Held and Ivan Habernal. 2025. [Contemporary LLMs struggle with extracting formal legal arguments](#). In *Proceedings of the Natural Language Processing Workshop 2025*, pages 292–303, Suzhou, China. Association for Computational Linguistics.
- Maya Ellen Hertz and Asta Sofie Stage Jarlner. 2025. [Trans “enough” for protection? experimenting with credibility in refugee status determination](#). *Frontiers in Human Dynamics*, Volume 7 - 2025.
- Anna Høgenhaug, Thomas Gammeltoft-Hansen, and Asta Sofie Stage Jarlner. 2023. [Nordic asylum practice in relation to religious conversion: Insights from denmark, norway and sweden](#). *Legal and Protection Policy Research Series*.
- Andong Hua, Kenan Tang, Chenhe Gu, Jindong Gu, Eric Wong, and Yao Qin. 2025. [Flaw or artifact? rethinking prompt sensitivity in evaluating LLMs](#). In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 19889–19899, Suzhou, China. Association for Computational Linguistics.
- IBM. 2025. [Ibm granite 4.0: hyper-efficient, high performance hybrid models for enterprise](#).
- Asta S Stage Jarlner, Maya Ellen Hertz, Marieke Anne Heyl, Thomas Gammeltoft-Hansen, and William Hamilton Byrne. 2026. [Credibility as a fuzzy concept in refugee law: a systematic literature review](#). *Journal of Ethnic and Migration Studies*, pages 1–31.
- Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. 2022. [Large language models are zero-shot reasoners](#). In *Proceedings of the 36th International Conference on Neural Information Processing Systems, NIPS ’22*, Red Hook, NY, USA. Curran Associates Inc.
- Viet Dac Lai, Nghia Ngo, Amir Pouran Ben Veyseh, Hieu Man, Franck Dernoncourt, Trung Bui, and Thien Huu Nguyen. 2023. [ChatGPT beyond English: Towards a comprehensive evaluation of large language models in multilingual learning](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 13171–13189, Singapore. Association for Computational Linguistics.
- Alexander H. Liu, Kartik Khandelwal, Sandeep Subramanian, Victor Jouault, Abhinav Rastogi, Adrien Sadé, Alan Jeffares, Albert Jiang, Alexandre Cahill, Alexandre Gavaudan, Alexandre Sablayrolles,

- Amélie Héliou, Amos You, Andy Ehrenberg, Andy Lo, Anton Eliseev, Antonia Calvi, Avinash Sooriyarachchi, Baptiste Bout, and 101 others. 2026. [Ministral 3](#). *Preprint*, arXiv:2601.08584.
- Laura Majer and Jan Šnajder. 2024. [Claim check-worthiness detection: How well do LLMs grasp annotation guidelines?](#) In *Proceedings of the Seventh Fact Extraction and VERification Workshop (FEVER)*, pages 245–263, Miami, Florida, USA. Association for Computational Linguistics.
- Taylor Meizlish and Chris Ziffo. 2026. [Evaluating an llm’s performance in annotating discourse strategies](#). *Corpus Pragmatics*, 10(1):23.
- Moran Mizrahi, Guy Kaplan, Dan Malkin, Rotem Dror, Dafna Shahaf, and Gabriel Stanovsky. 2024. [State of what art? a call for multi-prompt LLM evaluation](#). *Transactions of the Association for Computational Linguistics*, 12:933–949.
- Krzysztof Ociepa, Łukasz Flis, Remigiusz Kinas, Adrian Gwoździej, Krzysztof Wróbel, SpeakLeash Team, and Cyfronet Team. 2025a. [Bielik-11b-v3.0-instruct model card](#).
- Krzysztof Ociepa, Łukasz Flis, Remigiusz Kinas, Krzysztof Wróbel, and Adrian Gwoździej. 2025b. [Bielik 11b v3: Multilingual large language model for european languages](#). *Preprint*, arXiv:2601.11579.
- Maja Pavlovic and Massimo Poesio. 2024. [The effectiveness of LLMs as annotators: A comparative overview and empirical analysis of direct representation](#). In *Proceedings of the 3rd Workshop on Perspectivist Approaches to NLP (NLPerspectives) @ LREC-COLING 2024*, pages 100–110, Torino, Italia. ELRA and ICCL.
- Miguel Moura Ramos, Duarte M. Alves, Hippolyte Gisserot-Boukhlef, João Alves, Pedro Henrique Martins, Patrick Fernandes, José Pombal, Nuno M. Guerreiro, Ricardo Rei, Nicolas Boizard, Amin Farajian, Mateusz Klimaszewski, José G. C. de Souza, Barry Haddow, François Yvon, Pierre Colombo, Alexandra Birch, and André F. T. Martins. 2026. [Eurollm-22b: Technical report](#). *Preprint*, arXiv:2602.05879.
- Trine Rask Nielsen and Naja Holten Møller. 2022. [Data as a lens for understanding what constitutes credibility in asylum decision-making](#). *Proc. ACM Hum.-Comput. Interact.*, 6(GROUP).
- Mattes Ruckdeschel. 2025. [Just read the codebook! make use of quality codebooks in zero-shot classification of multilabel frame datasets](#). In *Proceedings of the 31st International Conference on Computational Linguistics*, pages 6317–6337, Abu Dhabi, UAE. Association for Computational Linguistics.
- Maximilian Schall and Gerard de Melo. 2025. [The hidden cost of structure: How constrained decoding affects language model performance](#). In *Proceedings of the 15th International Conference on Recent Advances in Natural Language Processing - Natural Language Processing in the Generative AI Era*, pages 1074–1084, Varna, Bulgaria. INCOMA Ltd., Shoumen, Bulgaria.
- Shivalika Singh, Angelika Romanou, Clémentine Fourrier, David Ifeoluwa Adelani, Jian Gang Ngui, Daniel Vila-Suero, Peerat Limkonchotiwat, Kelly Marchisio, Wei Qi Leong, Yosephine Susanto, Raymond Ng, Shayne Longpre, Sebastian Ruder, Wei-Yin Ko, Antoine Bosselut, Alice Oh, Andre Martins, Leshem Choshen, Daphne Ippolito, and 4 others. 2025. [Global MMLU: Understanding and addressing cultural and linguistic biases in multilingual evaluation](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 18761–18799, Vienna, Austria. Association for Computational Linguistics.
- Gemma Team, Aishwarya Kamath, Johan Ferret, Shreya Pathak, Nino Vieillard, Ramona Merhej, Sarah Perrin, Tatiana Matejovicova, Alexandre Ramé, Morgane Rivière, Louis Rouillard, Thomas Mesnard, Geoffrey Cideron, Jean bastien Grill, Sabela Ramos, Edouard Yvinec, Michelle Casbon, Etienne Pot, Ivo Penchev, and 197 others. 2025. [Gemma 3 technical report](#). *Preprint*, arXiv:2503.19786.
- Qwen Team. 2024. [Qwen2.5: A party of foundation models](#).
- Petter Törnberg. 2024. [Best practices for text annotation with large language models](#). *Preprint*, arXiv:2402.05129.
- Shubham Vatsal and Harsh Dubey. 2024. [A survey of prompt engineering methods in large language models for different nlp tasks](#). *Preprint*, arXiv:2407.12994.
- Yuqing Wang and Yun Zhao. 2024. [Metacognitive prompting improves understanding in large language models](#). In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 1914–1926, Mexico City, Mexico. Association for Computational Linguistics.
- Chuanchi Wen, Paul Clough, Rachel Paton, and Rebecca Middleton. 2025. [Leveraging large language models for thematic analysis: a case study in the charity sector](#). *AI & SOCIETY*, pages 1–18.
- Brandon T. Willard and Rémi Louf. 2023. [Efficient guided generation for large language models](#). *Preprint*, arXiv:2307.09702.
- Weihao Xuan, Rui Yang, Heli Qi, Qingcheng Zeng, Yunze Xiao, Aosong Feng, Dairui Liu, Yun Xing, Junjue Wang, Fan Gao, Jinghui Lu, Yuang Jiang, Huitao Li, Xin Li, Kunyu Yu, Ruihai Dong, Shangding Gu, Yuekang Li, Xiaofei Xie, and 13 others. 2025. [MMLU-ProX: A multilingual benchmark for advanced large language model evaluation](#). In *Proceedings of the 2025 Conference on Empirical Methods*

in *Natural Language Processing*, pages 1513–1532, Suzhou, China. Association for Computational Linguistics.

Mike Zhang, Max Müller-Eberstein, Elisa Bagnana, and Rob van der Goot. 2024. [Snakmodel: Lessons learned from training an open danish large language model](#). *Preprint*, arXiv:2412.12956.

Jingming Zhuo, Songyang Zhang, Xinyu Fang, Haodong Duan, Dahua Lin, and Kai Chen. 2024. [ProSA: Assessing and understanding the prompt sensitivity of LLMs](#). In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 1950–1976, Miami, Florida, USA. Association for Computational Linguistics.

Caleb Ziems, William Held, Omar Shaikh, Jiaao Chen, Zhehao Zhang, and Diyi Yang. 2024. [Can large language models transform computational social science?](#) *Computational Linguistics*, 50(1):237–291.

## A Dataset details

Here we present details about how the RAB-Cred dataset was sampled and annotated, and how we extract the outcome for each case.

### A.1 Intended use

We provide multi-annotator labels and confidence levels for RAB written decisions indicating the presence and sentiment of credibility assessment, as defined by domain experts, along with case metadata (e.g. outcome). These annotations and metadata must only be used for research purposes, and are solely intended to be used for understanding *whether* and *how* credibility assessment is performed in asylum decisions. They are **not** intended to be used for assessing/classifying the credibility or veracity of claims made by the applicant. Furthermore, they are **not** intended to be used for automating decisions.

### A.2 Dataset source

The dataset was collected via web-scraping of all written decisions available at <https://fln.dk/praksis/> in June 2025 and early December 2025. Combined, this yielded 10817 unique cases dating from 2004 to 2025, from which we sampled two subsets to annotate (val and test).

**Representativeness** Not all RAB decisions are published on the website. The website states: "The Refugee Appeals Board's website regularly publishes summaries of selected decisions that represent the Board's practice regarding individual countries. This means that not all of the Board's decisions are published on the website."<sup>1</sup> (translated with DeepL.com). The RAB's specific selection criteria is unknown.

Empirically, we compared the recognition rate (outcome) and country of origin distribution of the scraped data to the yearly statistics published by the RAB<sup>2</sup>. Overall, before 2015 we have a slight over-representation of over-turned cases. In terms of country of origin, the scraped data also resembles the top-3 national distribution reported by the RAB, but with a slight under-representation of soviet states and over-representation of Middle East cases.

**Anonymization** Written decisions published by the RAB are pseudo-anonymized. Although the specific pseudo-anonymization procedure and criteria is unknown and appears to have evolved over time, the following is stated on the RAB's website: "The practice summaries reproduce the Refugee Board's reasoning in each individual decision in full. However, in some cases, names, dates, locations, etc. have been anonymised for the sake of the applicant." (translated with DeepL.com)<sup>3</sup>.

In practice, we observe that no names or initials of applicants or their relatives are present in any case texts. Furthermore, in cases from 2010 or later, specific details (e.g. a date, age, country, city, ethnicity, medical issue, social media platform, among others) are redacted in the written decisions, and replaced by square brackets.

### A.3 Sampling of the validation and test sets

The validation set and test set were sampled separately from the 10817 web-scraped cases. Sampling was performed by the domain experts H1 and H2. Yearly distribution for the validation and test sets can be seen in Figure A.1 and is detailed below.

**Validation set** Cases were randomly sampled with yearly stratification across the year range 2004 to 2021 (18 years), with 4 cases per year. In addition, 1 recent case (from 2025) was added due to its difficult nature (surplace case, multiple asylum motives). This results in 73 cases.

<sup>1</sup>"På Flygtningenævnets hjemmeside offentliggøres løbende praksisresumeer af udvalgte afgørelser, der udgør et repræsentativt udsnit af nævnets praksis vedrørende de enkelte lande. Det er således ikke alle nævnets afgørelser, der offentliggøres på hjemmeside." [https://web.archive.org/web/20260302132511/https://fln.dk/information\\_til/advokater/naevnets\\_praksis/](https://web.archive.org/web/20260302132511/https://fln.dk/information_til/advokater/naevnets_praksis/)

<sup>2</sup>[https://fln.dk/statistik\\_og\\_maaltal/](https://fln.dk/statistik_og_maaltal/)

<sup>3</sup>"I praksisresumeerne gengives Flygtningenævnets præmis i den enkelte afgørelse i sin fulde længde. Der er dog i nogle tilfælde af hensyn til ansøgeren foretaget en anonymisering af navne, tidsangivelser, stedsangivelser etc." [https://web.archive.org/web/20260302132511/https://fln.dk/information\\_til/advokater/naevnets\\_praksis/](https://web.archive.org/web/20260302132511/https://fln.dk/information_til/advokater/naevnets_praksis/)

**Test set** Cases cover a range of years between 2004 to 2025, where time ranges for stratification are based on changes to the board: 2004-2012, 2013-2016, 2017-2021, 2022-2025. The board changes is in its size and composition, and thus could be affected by slight changes in writing style and possibly practice.

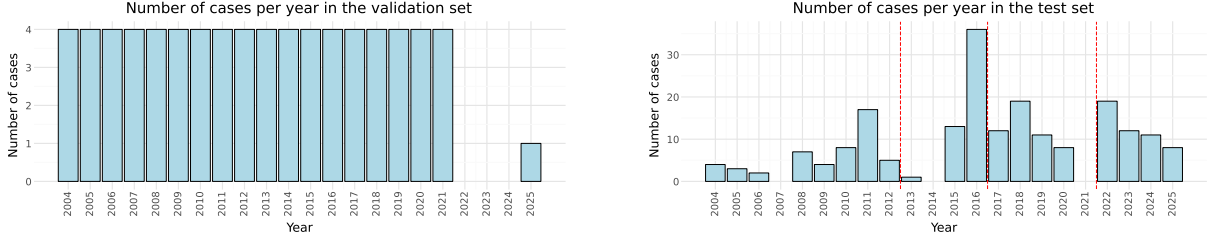


Figure A.1: Number of cases in the RAB-Cred dataset by year. The red lines show the 4 time ranges used for stratified sampling of the test set.

**A.4 Human annotators**

Annotation of the validation and test set was performed by the same two annotators (H1 and H2). A third annotator (H3) was introduced to resolve the 4 test set samples where H1 and H2 assigned conflicting labels.

All three annotators are fluent Danish speakers (H2 and H3 being native speakers), with a background in social science and several years of research experience in the field of Danish refugee law, and highly familiar with both the context and the content of the RAB decision texts.

**A.5 Test set annotations**

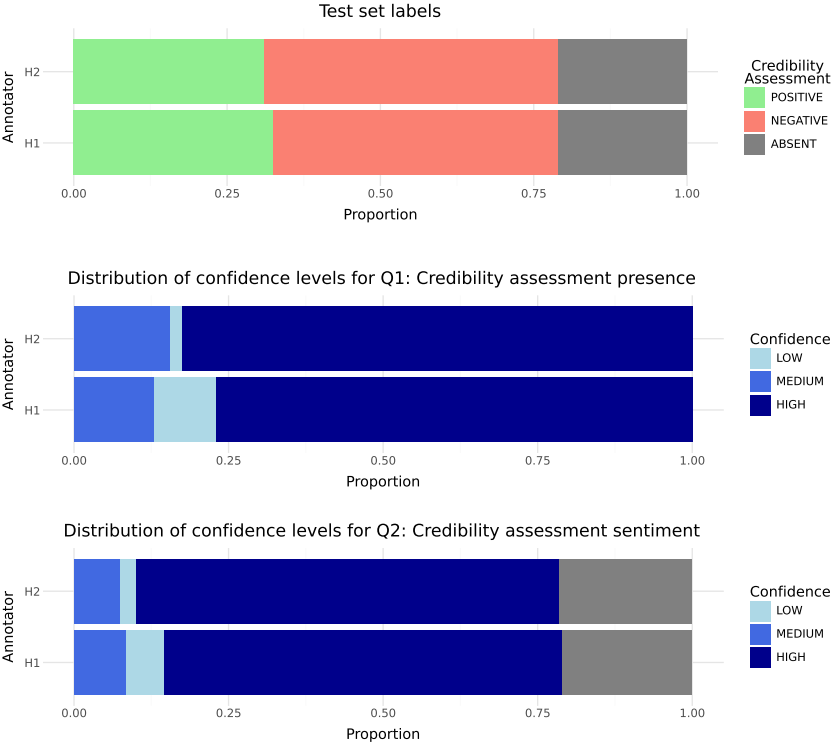


Figure A.2: Distribution of labels independently assigned by the two domain experts on the test set (top) and their confidence level for the presence and (optionally) sentiment of a credibility assessment.

## A.6 Outcome extraction

We apply regex-based pattern matching to extract the outcome of each case (rejection upheld, rejection reversed, or remanded). For the few cases (three across RAB-Cred) whose outcome was not automatically determined via pattern matching, we labeled the outcome manually.

```
MONTHS_DA = r'(januar|februar|marts|april|maj|juni|juli|august|september|oktober|november|december)'  
YEAR = r'\d{4}'  
IMMIGRATION_SERVICE = r'Udl nding(?:estyrelsens|eservice|eservices|estyrelsen)'  
  
rejection_upheld_patterns = [  
  rf'Flertallet stemte derfor for at tiltr de Udl ndingestyrelsens afg relse',  
  rf'N vnet stadf stede i {MONTHS_DA} {YEAR} {IMMIGRATION_SERVICE} afg relse',  
  rf'Flygtningen vnet stadf ster derfor {IMMIGRATION_SERVICE} afg relse',  
  rf'staf ster Flygtningen vnet derfor {IMMIGRATION_SERVICE} afg rels',  
  rf'N vnet stadf stede i {MONTHS_DA}.*{YEAR} {IMMIGRATION_SERVICE} afg relse',  
  rf'ikke betingelserne for opholdstilladelse',  
  rf'ikke sandsynliggjort, at ans geren.*?vil risikere forf lgelse',  
  rf'fandt.*?ikke, at ans geren havde krav p  opholdstilladelse',  
  rf'ikke, at ans gerens.*?ville v re i en s dan risiko herfor, at der var grundlag for at meddele asyl',  
  rf'ikke.*?at ans gerne skulle meddeles opholdstilladelse',  
  rf'kan det ikke antages, at ans geren ved en tilbagevenden skulle v re i en reel risiko',  
  rf'Flygtningen vnets afslag',  
  rf'ikke.*?at den kan begrunde opholdstilladelse efter udl ndingelovens',  
  rf'Flygtningen vnet finder.*?ikke, at ans geren.*?risikerer',  
  rf'ikke, at ans geren.*?ville v re i risiko',  
  rf'ikke fandtes at v re asylbegrundende',  
  rf'Flygtningen vnet fandt.*?ikke, at ans geren ved en tilbagevenden.*?ville risikere',  
  rf'ikke, at ans geren opfylder betingelserne for asyl',  
  rf'opfyldte ans geren ikke betingelserne for at f  asyl',  
  rf'han ikke kunne p ber be sig den beskyttelse, som f lger af udl ndingelovens',  
  rf'Flygtningen vnet fandt ikke, at disse forhold kunne begrunde',  
  rf'meddeler derfor ans geren afslag p  opholdstilladelse',  
  rf'staf ster Flygtningen vnet {IMMIGRATION_SERVICE} afg relse',  
  rf'staf ster herefter {IMMIGRATION_SERVICE} afg relse',  
  rf'staf stede i (?:{MONTHS_DA}) {YEAR} {IMMIGRATION_SERVICE} afg relse',  
  rf'finder Flygtningen vnet heller ikke,? at det vil v re uproportionalt.*?opholdstilladelse'  
]  
  
rejection_reversed_patterns = [  
  rf'Klageren opfylder s ledes betingelserne for at blive meddelt opholdstilladelse',  
  rf'N vnet meddelte i opholdstilladelse(.*)til',  
  rf'N vnet meddelte i (?:{MONTHS_DA}) {YEAR} opholdstilladelse',  
  rf'Flygtningen vnet oph vede derfor {IMMIGRATION_SERVICE}',  
  rf'Flygtningen vnet  ndrer derfor {IMMIGRATION_SERVICE} afg relse',  
  rf'finder Flygtningen vnet s ledes, at De skal meddeles opholdstilladelse',  
  rf'N vnet omgjorde i (?:{MONTHS_DA}) {YEAR} {IMMIGRATION_SERVICE} afg relse',  
  rf'N vnet genoptog og omgjorde i (?:{MONTHS_DA}) {YEAR}',  
  rf'besluttet at genoptage sagen og omg re {IMMIGRATION_SERVICE} afg relse',  
  rf'meddeler derfor klageren opholdstilladelse',  
  rf'N vnet  ndrede i (?:{MONTHS_DA}) {YEAR} {IMMIGRATION_SERVICE} afg relse',  
  rf'klageren meddeles opholdstilladelse efter udl ndingelovens'  
]  
  
remanded_patterns = [  
  rf'N vnet hjemviste i',  
  rf'sagen b r hjemvises til {IMMIGRATION_SERVICE}',  
  rf'N vnet hjemviste i (?:{MONTHS_DA}) {YEAR} {IMMIGRATION_SERVICE} afg relse'  
]
```

Figure A.3: Regex patterns used to label the outcome of each case.

## B Experimental set-up

Here we present the prompt variants and models used in our experiments, along with implementation details for text classification.

### B.1 Prompt variants

#### B.1.1 System prompts

System prompts follow a nested structure, where SP2, SP3, SP4 and SP5 extend SP1. Furthermore, SP4 extends SP3 with edge cases.

(SP0) No system prompt

(SP1) Persona

You are a highly experienced legal scholar and evidence analyst specializing in qualitative coding of Danish asylum cases. You have in-depth knowledge and understanding of asylum decision-making in Denmark and the crucial role of risk and credibility. You are thorough, analytical and able to dissect complex legal reasonings.

You will analyze asylum decision summaries to identify the presence and the sentiment of a credibility assessment.

(SP2) Persona + Codebook

**(SP1)**

## Context: What is a credibility assessment?

The purpose of credibility assessment is to determine if the claimants account can be accepted as true (in whole or in part) or not. The focus is on the past and present facts asserted by the claimant. The credibility assessment relates to establishing facts, not whether these facts justify protection or whether there is a risk of persecution.

Notably, it is different from the risk assessment, which is conducted to determine if there, based on the established facts, is a risk of persecution or ill treatment, reaching the legal threshold for protection. This assessment is forward looking considering what would happen if the claimant was returned.

Code no credibility assessment in the absence of a credibility assessment.

## Context: Credibility sentiment:

The overall credibility sentiment captures the decision-maker's overall conclusion on whether the information provided by the claimant can be accepted as facts or not. It reflects the net outcome, not the presence of isolated inconsistencies or doubts.

Code credibility as overall positive where a credibility assessment is present and the decision-maker accepts the core of the claimant's account as true, even if some details are questioned, discounted, or left unresolved. The risk assessment can proceed based on the information provided by the claimant, also when there are minor doubts. The acceptance of facts may be argued or it may simply be stated that the Board accepts the motive and then proceed to a risk assessment. Please distinguish between when the board accepts the facts and proceed to risk assessment and when the board notes that even if the board accepted the facts, the risk assessment would be negative.

Code credibility as overall negative where a credibility assessment is present and where the decision-maker rejects the core of the claimant's account, such that the presented factual basis of the claim is not accepted as factual. Some facts may be accepted, but the central aspects motivating protection are rejected. The risk assessment proceeds without taking the central elements of the information provided by the claimant into the account, even if parts of claim are accepted.

(SP3) Persona + Computer scientist prompt

(SP1)

## Context: What is a credibility assessment?

Credibility Assessment: Evaluation of whether the claimant's factual account can be accepted. This establishes facts about past/present events. This is BACKWARD-LOOKING (what happened).

NOT Credibility Assessment:

- Risk assessment: Forward-looking evaluation (will persecution occur if returned?)
- Legal threshold analysis: Whether proven facts meet asylum criteria
- Explicitly bypassing credibility assessment: Even if the claimant's explanation were accepted, it would not qualify for asylum.

## Context: Classification categories / codes

"NO CREDIBILITY ASSESSMENT" - Credibility not assessed

- The decision proceeds directly to risk/legal analysis
- Or there are references to credibility without establishing the credibility sentiment, as the credibility of presented facts is not considered relevant for the asylum decision
- Danish phrases indicating this: "uanset om... lægges til grund "

"POSITIVE CREDIBILITY ASSESSMENT" - Credibility assessed, core account ACCEPTED as true

- Risk assessment proceeds WITH the claimant's information
- Minor doubts about peripheral details don't change this
- Core of the claimant's account accepted even if some inconsistencies noted
- Danish phrases: "forklaring(en) ... lægges til grund", "ansøgers forklaring er troværdig"

"NEGATIVE CREDIBILITY ASSESSMENT" - Credibility assessed, core account REJECTED

- Risk assessment proceeds WITHOUT central elements of claimant's account
- Core of the claimant's account not believed, even if some peripheral facts accepted
- Danish phrases: "forklaring(en) ... ikke lægges til grund", "ikke troværdig", "divergerende", "konstrueret", "usandsynlig"

IMPORTANT NOTE: Case outcome is different from Credibility assessment

- Asylum claim can be REJECTED despite facts being believed (positive credibility assessment)
- Asylum claim can meet the legal threshold for asylum and thus be ACCEPTED despite the claimant's account being deemed non-credible (negative credibility assessment).

(SP4) Persona + Computer scientist prompt

(SP3)

## Context: critical edge cases

1. Hypothetical Construction = "NO CREDIBILITY ASSESSMENT"

- "Even if X was accepted, it wouldn't meet the legal threshold of protection..."
- This is NOT evaluating if X is true, just analyzing consequences IF it were true

2. Minor Doubts + Core Accepted as True = "POSITIVE CREDIBILITY ASSESSMENT"

- "Some inconsistencies about dates, but the political involvement is accepted..."
- The core claim is believed despite peripheral doubts

3. Some Facts Accepted + Core Not Believed = "NEGATIVE CREDIBILITY ASSESSMENT"

- "X is accepted, but not the claimed persecution..."
- Peripheral facts accepted but asylum grounds rejected

(SP5) Persona + Domain expert prompt

(SP1)

BACKGROUND: Credibility assessment (troværdighedsvurdering) refers to evaluations of whether the asylum seeker's account is believable.

## Context: CRITICAL DISTINCTIONS:

NEUTRAL REPORTING vs. CREDIBILITY ASSESSMENT e.g.,:

- NOT credibility: 'Ansøgeren har henvist til...' (The applicant has referred to...)
- NOT credibility: 'Ansøgeren har forklaret...' (The applicant has explained...)
- IS credibility: 'Flygtningenævnet lægger i det væsentlige ansøgerens forklaring til grund' OR 'Flygtningenævnet lægger ikke ansøgerens forklaring til grund'
- IS credibility: 'Nævnet finder ansøgerens forklaring troværdig' (evaluation)
- IS credibility: 'Ansøgerens forklaring er ikke sammenhængende' (consistency assessment)
- IS credibility: 'Ansøgers forklaring fremstår selv-oplevet og detaljeret' OR 'ansøgers forklaring fremstår ikke selv-oplevet og detaljeret' (detail assessment and genuineness assessment)

## Context: LEGAL ASSESSMENT vs. CREDIBILITY ASSESSMENT:

THE KEY DISTINCTION: Legal assessments evaluate whether PROVEN FACTS meet legal criteria. Credibility assessments evaluate whether the applicant's TESTIMONY is believable.

- NOT credibility: Legal outcome formulations, evaluating whether events meet legal thresholds for asylum
- NOT credibility: Assessments of severity, intensity, or legal thresholds (e.g., 'disse forhold har haft et sådan omfang og intensitet')
- NOT credibility: Risk assessments based on country conditions or objective circumstances
- IS credibility: Evaluating whether the applicant's ACCOUNT of events is believable
- IS credibility: Assessing the truthfulness or consistency of the applicant's TESTIMONY

## Context: CORE PRINCIPLE: When analyzing a case, ask yourself the key question: 'Is the decision-maker assessing whether something happened, and/or the applicant's fear/convictions/political opinions/sexual orientation are genuine [IS CREDIBILITY], OR are they evaluating the risk and legal implications of accepted facts? [NOT CREDIBILITY]'

## B.1.2 User prompts

(UP1)

Analyze the following legal case decision from the Flygtningenævnet (Danish Refugee Appeals Board) and classify it into one of the following three categories: [NO CREDIBILITY ASSESSMENT, POSITIVE CREDIBILITY ASSESSMENT, NEGATIVE CREDIBILITY ASSESSMENT]. Your answer should only contain the chosen category in capital letters.

Case text: [INSERT CASE TEXT HERE]

Category:

(UP2)

Analyze the following legal case decision from the Flygtningenævnet (Danish Refugee Appeals Board) and determine whether credibility is assessed.

Case text: **[INSERT CASE TEXT HERE]**

Is there a credibility assessment? Your answer should only contain "Y" or "N".

---

*If the answer is "N", stop here.*

*If the answer is "Y", proceed to the follow-up prompt:*

---

Now, determine the sentiment of the credibility assessment. Is the credibility assessment POSITIVE or NEGATIVE? Your answer should only contain the chosen category in capital letters.

(UP1-FS) Few-shot

Analyze the following legal case decision from the Flygtningenævnet (Danish Refugee Appeals Board) and classify it into one of the following three categories: [NO CREDIBILITY ASSESSMENT, POSITIVE CREDIBILITY ASSESSMENT, NEGATIVE CREDIBILITY ASSESSMENT]. Your answer should only contain the chosen category in capital letters.

Case text: **[INSERT EXAMPLE 1 HERE]**

Category: NO CREDIBILITY ASSESSMENT

Case text: **[INSERT EXAMPLE 2 HERE]**

Category: POSITIVE CREDIBILITY ASSESSMENT

Case text: **[INSERT EXAMPLE 3 HERE]**

Category: NEGATIVE CREDIBILITY ASSESSMENT

Case text: **[INSERT CASE TEXT HERE]**

Category:

(UP3) Zero-shot Chain-of-Thought

Analyze the following legal case decision from the Flygtningenævnet (Danish Refugee Appeals Board) and classify it into one of the following three categories: [NO CREDIBILITY ASSESSMENT, POSITIVE CREDIBILITY ASSESSMENT, NEGATIVE CREDIBILITY ASSESSMENT].

Case text: **[INSERT CASE TEXT HERE]**

Let's think step-by-step.

---

To conclude, classify the case into one of the following three categories: [NO CREDIBILITY ASSESSMENT, POSITIVE CREDIBILITY ASSESSMENT, NEGATIVE CREDIBILITY ASSESSMENT]. Your answer should only contain the chosen category in capital letters.

#### (UP4) Zero-shot Metacognitive

Analyze the following legal case decision from the Flygtningenævnet (Danish Refugee Appeals Board) and classify it into one of the following three categories: [NO CREDIBILITY ASSESSMENT, POSITIVE CREDIBILITY ASSESSMENT, NEGATIVE CREDIBILITY ASSESSMENT].

Case text: [INSERT CASE TEXT HERE]

As you perform this task, follow these steps:

1. Clarify your understanding of the task and the case.
2. Make a preliminary judgment on whether a credibility assessment is made in the case, and if so, its sentiment.
3. Critically assess your preliminary analysis. If you are unsure about your initial judgment, reassess it.
4. Confirm your final answer and explain the reasoning behind your decision.
5. Evaluate your confidence (0-100%) in your analysis and provide an explanation for this confidence level.

---

To conclude, classify the case into one of the following three categories: [NO CREDIBILITY ASSESSMENT, POSITIVE CREDIBILITY ASSESSMENT, NEGATIVE CREDIBILITY ASSESSMENT]. Your answer should only contain the chosen category in capital letters.

**Few-shot examples** The following validation set cases are used as few-shot examples:

- EXAMPLE 1: No credibility assessment: ID 19715 <https://fln.dk/praksis/2020/december/syri202051/>
- EXAMPLE 2: Positive credibility assessment: ID 18571 <https://fln.dk/praksis/2019/november/egyp20194/>
- EXAMPLE 3: Negative credibility assessment: ID 20194 <https://fln.dk/praksis/2021/februar/syri202118/>

#### Example 1: No Credibility Assessment (Case text translated with DeepL.com)

In August 2020, the Board upheld the Danish Immigration Service's decision regarding a female citizen of Syria. She arrived in 2016. The Refugee Board stated: "The complainant is an ethnic Arab from Hajar Al-Aswad, Damascus, Syria. The appellant has been politically active on Facebook, where she shares critical posts about conditions in Syria. The case file shows that the appellant entered Denmark in [the summer] of 2016 and that she was granted a residence permit in [the fall] of 2019 pursuant to Section 7(3) of the Aliens Act. In the fall of 2020, the complainant appealed to the Refugee Board against the Danish Immigration Service's decision, claiming a right to a residence permit pursuant to Section 7(1) of the Aliens Act. In support of this, the complainant has stated that she fears the Syrian authorities will subject her to violence, kidnapping, and rape because her family is wanted. The complainant has stated that her maternal uncles have been arrested, killed, or forced to flee to Jordan because one of her maternal uncles had smashed the president's statue, and photos of him with his foot on the statue have been circulated. After the Syrian authorities burned the complainant's maternal uncles' houses and motorcycles, the complainant fled with her mother and sisters to Jordan. The complainant's maternal aunts have been arrested and are being punished by the Syrian regime. The complainant's father and several of her father's family members have been granted residence permits under Section 7(1) of the Aliens Act. In support of her asylum claim, the complainant has further stated that she fears the Syrian authorities because she is politically active on Facebook. The complainant has stated in this regard that, should she return, she fears the war and the bombings. The complainant has stated that she fears the government, the president, and the military. The complainant has stated that she uses Facebook to share information about the war. The complainant primarily writes the posts she shares herself. The complainant primarily shares posts with political content, as they concern the situation in Syria. The complainant's Facebook profile is public, as she wants everyone to know what is happening. The complainant began sharing posts about Syria in late 2018. In the event of a return to Syria, the complainant fears that the authorities will search her phone and arrest her based on her Facebook activism. The Refugee Board finds—like the Danish Immigration Service—that it is based on the complainant's own assumption that her Facebook posts, which according to her are critical of the regime, have come to the attention of the Syrian authorities. In this regard, the Refugee Board has emphasized that the complainant has a limited number of "friends" on Facebook, and that these consist solely of friends and family. Furthermore, and since the complainant—contrary to what

was stated in a concurrent decision in her mother’s case— be at risk due to her kinship with her uncle, the Refugee Board does not find that the appellant has established a likelihood that, upon return to Syria, she would be at risk of persecution or abuse covered by Section 7(1) or (2) of the Aliens Act. The appellant thus does not meet the conditions for being granted a residence permit pursuant to Section 7(1) or Section 7(2) of the Aliens Act, and therefore the Refugee Board upholds the Danish Immigration Service’s decision.” Syri/2020/51/EDO

**Annotators H1 and H2:** Q1: Credibility assessment present? **No** (Confidence: High)

**Outcome:** Rejection upheld

### Example 2: Positive Credibility Assessment (Case text translated with DeepL.com )

In October 2019, the Board upheld the Danish Immigration Service’s decision regarding a male citizen of Egypt. He entered the country in 2016. The Refugee Board stated: “The applicant is an ethnic Arab and a Coptic Christian from [town name], Dakahliya, Egypt. The applicant has not been a member of political or religious associations or organizations, nor has he otherwise been politically active. As grounds for asylum, the applicant has cited his fear of being killed by his own family or random Egyptians because he is homosexual. In support of this, he has stated that it is no longer a secret that he is homosexual, as people living here have become aware of this. The applicant’s father has learned of the applicant’s sexuality through these people and has threatened over the phone that he will cut the applicant into small pieces and have nothing to do with him. He has also cited as a basis for his asylum claim that he fears being killed by two men named [A] and [B] or by random individuals because he is Christian. In support of this, he has stated that at one point he picked up [A] and [B] in a tuk-tuk and drove them to a deserted area. [B] pulled out a knife, and the applicant assumed he intended to kill him, so he fled the scene. The applicant reported the incident to the police, who arrested [A] and [B]. During the trial in 2005, the applicant was offered money by [A]’s family to withdraw the complaint. He received an indirect threat from [A]’s maternal grandfather. The applicant subsequently dropped the case against [A] and [B], as he did not believe he would gain anything from their imprisonment. The applicant has had no problems with them since 2005. However, the applicant suspects that there is a connection between the case against [A] and [B] and the fact that the applicant’s father was wrongfully accused of various criminal offenses from 2005 to 2010. From 2006 until his departure on [date in the spring] 2016, the applicant lived in both Cairo and Sharm el-Sheikh without experiencing any problems with [A] or [B]. He has further noted that random attacks on Coptic Christians occur in Egypt. Finally, as a basis for his asylum claim, the applicant has stated that he fears the Egyptian intelligence service because he has applied for asylum in Denmark. The Refugee Board, like the Danish Immigration Service, finds the applicant’s explanation credible. However, the Refugee Board finds, for the same reasons stated by the Danish Immigration Service, that neither the general conditions for Coptic Christians in Egypt, the applicant’s fear of the Egyptian intelligence service, nor the applicant’s conflict with [A], [B], and their families can justify a residence permit under Section 7 of the Aliens Act. It appears from the available background information that it is not illegal to be an LGBT person or to have sex with a person of the same sex in Egypt. Notwithstanding that the available background information also indicates that in September 2017 the authorities adopted a tougher stance toward homosexuals and that, in connection with this, a bill was introduced to criminalize homosexuality, the Refugee Board does not find that the general conditions —despite being difficult—for homosexuals in Egypt can in themselves lead to the applicant being granted a residence permit. Nor does the majority of the Refugee Board find that the applicant has substantiated that he has been specifically and individually persecuted as a result of his sexuality. In this regard, the majority has emphasized that, based on the information regarding the telephone threat from the applicant’s father, it has not been substantiated that the father will actively seek out the applicant and carry out his threat. The majority thus finds that the applicant can be directed to reside outside his father’s home region, including in Cairo, where he has been able to reside without issue for approximately 4 years starting in 2005. The fact that, according to the applicant’s account, other individuals in Egypt have become aware of his sexuality cannot lead to a different assessment. The majority of the Refugee Board thus finds, after a comprehensive assessment, that the applicant has not established that he would be at risk of persecution or abuse covered by Section 7 of the Aliens Act upon a return to Egypt. The Refugee Board therefore upholds the decision of the Danish Immigration Service.” Eyp/2019/4/DH

**Annotators H1 and H2:** Q1: Credibility assessment present? **Yes** (Confidence: High)

**Annotators H1 and H2:** Q2: Credibility assessment sentiment **Positive** (Confidence: High)

**Outcome:** Rejection upheld

### Example 3: Negative Credibility Assessment (Case text translated with DeepL.com )

In February 2021, the Board overturned the Danish Immigration Service’s decision regarding a female citizen of Syria. She entered the country in 2016. The Refugee Board stated: “The appellant is an ethnic Kurd and a Sunni Muslim from Dirik, Al Hasakah, Syria. The complainant has not been a member of any political or religious associations or organizations, nor has she been politically active in any other way. The Danish Immigration Service granted the appellant a residence permit in the summer of 2016 pursuant to Section 7(3) of the Aliens Act. During the original asylum case, the appellant cited as her reason for seeking asylum that she feared the general conditions in Syria if she were to return. In support of

this, the complainant had stated that there was a war in Syria and that the regular bombings in Syria were affecting her mental state. The complainant had further stated that she wished to live a peaceful life with her children. In the summer of 2020, the Danish Immigration Service decided to refuse to extend the complainant's residence permit pursuant to Section 11(2), second sentence, of the Aliens Act. The Danish Immigration Service has assessed that the basis for the complainant's residence permit no longer exists. The complainant has continued to cite general conditions as grounds for asylum, as well as her family's political activities and conflicts with the authorities. For the reasons stated by the Danish Immigration Service, the Refugee Board agrees that the complainant has not demonstrated that, upon return to Syria, she would be at risk of persecution or abuse covered by Section 7(1) or (2) of the Aliens Act. The complainant's statements to the Refugee Board do not warrant a different assessment. In this regard, the Refugee Board has emphasized that the complainant's statements cannot, in essence, be relied upon, as the complainant has, among other things, provided detailed accounts of her spouse's, sons', and sons-in-law's political activities, and that, according to a neighbor, the family's home was visited by the authorities with the intention of arresting the family shortly before their departure. The fact that the complainant and family members in Denmark participated in a demonstration cannot lead to a different assessment either, as the complainant cannot be presumed to have come under the scrutiny of the Syrian authorities for that reason, regardless of whether photos of the demonstration may have been posted on Facebook. The Refugee Board further agrees with the Danish Immigration Service that the general conditions in Damascus are no longer such that anyone would be at a real risk of being subjected to abuse in violation of Article 3 of the European Convention on Human Rights solely as a result of their mere presence in the area. As stated by the Danish Immigration Service, the Refugee Board must now determine whether refusing to extend the complainant's residence permit would violate Article 8 of the European Convention on Human Rights. The complainant is Kurdish and 63 years old. She left Syria around 2012 and entered Denmark in 2016. The complainant is a single woman, illiterate, and has never been part of the labor market. The complainant has nine adult children, one of whom lives in Sweden, while the rest live in Denmark. The complainant has never lived alone and, since her arrival in Denmark, has lived with one of her adult sons and his family; the complainant's youngest son, aged 22, now also lives with her. Furthermore, based on the complainant's statement, the Refugee Board finds that she has absolutely no family or other network in Syria. On that basis and following a comprehensive and concrete assessment, the Refugee Board finds that it would currently be contrary to Denmark's international obligations, cf. Article 8 of the European Convention on Human Rights, to refuse to extend the complainant's residence permit. The Refugee Board therefore extends the complainant's residence permit pursuant to Section 7(3) of the Aliens Act." Syri/2021/18/HZC

---

**Annotators H1 and H2:** Q1: Credibility assessment present? **Yes** (Confidence: High)

**Annotators H1 and H2:** Q2: Credibility assessment sentiment **Negative** (Confidence: High)

**Outcome:** Rejection reversed

## B.2 Model selection and implementation

### B.2.1 Model selection

We evaluate the following 21 models on the validation set. The top-5 models which are selected for evaluation on the test set are in bold. All models are pulled from HuggingFace and used in their default precision.

	huggingface model	context length	multilingual capabilities mentioned in Hugging Face model card
Aya	<a href="#">CohereLabs/aya-expense-8b</a>	8K	"Languages covered: The model is particularly optimized for multilinguality and supports the following languages: Arabic, Chinese (simplified & traditional), Czech, Dutch, English, French, German, Greek, Hebrew, Hindi, Indonesian, Italian, Japanese, Korean, Persian, Polish, Portuguese, Romanian, Russian, Spanish, Turkish, Ukrainian, and Vietnamese"
	<a href="#">CohereLabs/aya-expense-32b</a>	128K	
Gemma	<b><a href="#">gemma-3-27b-it</a></b>	128K	"multilingual support in over 140 languages" "The training dataset includes content in over 140 languages."
	<b><a href="#">gemma-3-12b-it</a></b>		
	<b><a href="#">gemma-3-4b-it</a></b>		
Granite	<a href="#">ibm-granite/granite-4.0-micro</a>	128K	"Supported Languages: English, German, Spanish, French, Japanese, Portuguese, Arabic, Czech, Italian, Korean, Dutch, and Chinese. Users may finetune Granite 4.0 models for languages beyond these languages."
Llama	<a href="#">meta-llama/Llama-3.1-8B-Instruct</a>	128K	"Supported languages: English, German, French, Italian, Portuguese, Hindi, Spanish, and Thai." "Note: Llama 3.1 has been trained on a broader collection of languages than the 8 supported languages. Developers may fine-tune Llama 3.1 models for languages beyond the 8 supported languages "
	<a href="#">meta-llama/Llama-3.2-3B-Instruct</a>	128K	"Supported Languages: English, German, French, Italian, Portuguese, Hindi, Spanish, and Thai are officially supported. Llama 3.2 has been trained on a broader collection of languages than these 8 supported languages. Developers may fine-tune Llama 3.2 models for languages beyond these supported languages"
Phi	<b><a href="#">microsoft/phi-4</a></b>	16K	"Multilingual data constitutes about 8% of our overall data. " "The model is trained primarily on English text. Languages other than English will experience worse performance." "phi-4 is not intended to support multilingual use. "
	<a href="#">microsoft/Phi-4-mini-instruct</a>	128K	"Supported languages: Arabic, Chinese, Czech, Danish, Dutch, English, Finnish, French, German, Hebrew, Hungarian, Italian, Japanese, Korean, Norwegian, Polish, Portuguese, Russian, Spanish, Swedish, Thai, Turkish, Ukrainian" "The model is intended for broad multilingual commercial and research use." "The Phi models are trained primarily on English text and some additional multilingual text. Languages other than English will experience worse performance as well as performance disparities across non-English."
Mistral	<a href="#">mistralai/Mistral-3-8B-Instruct-2512</a>	256K	"Supports dozens of languages, including English, French, Spanish, German, Italian, Portuguese, Dutch, Chinese, Japanese, Korean, Arabic."
	<b><a href="#">mistralai/Mistral-3-14B-Instruct-2512</a></b>	128K	"Supports dozens of languages, including English, French, German, Greek, Hindi, Indonesian, Italian, Japanese, Korean, Malay, Nepali, Polish, Portuguese, Romanian, Russian, Serbian, Spanish, Swedish, Turkish, Ukrainian, Vietnamese, Arabic, Bengali, Chinese, Farsi."
Qwen	<a href="#">Qwen/Qwen2.5-32B-Instruct</a>	256K	"Multilingual support for over 29 languages, including Chinese, English, French, Spanish, Portuguese, German, Italian, Russian, Japanese, Korean, Vietnamese, Thai, Arabic, and more."
	<a href="#">Qwen/Qwen2.5-7B-Instruct</a>	131K	
	<a href="#">Qwen/Qwen2.5-3B-Instruct</a>	128K	
	<a href="#">Qwen/Qwen2.5-14B-Instruct</a>	128K	
	<b><a href="#">Qwen/Qwen3-30B-A3B-Instruct-2507</a></b>	256K	"Substantial gains in long-tail knowledge coverage across multiple languages."
	<a href="#">Qwen/Qwen3-4B-Instruct-2507</a>		
Bielik	<a href="#">speakleash/Bielik-11B-v3.0-Instruct</a>	32K	"Developed and trained on multilingual text corpora across 32 European languages, with emphasis on Polish"
EuroLLM	<a href="#">utter-project/EuroLLM-22B-Instruct-2512</a>	32K	"Language(s) (NLP): Bulgarian, Croatian, Czech, Danish, Dutch, English, Estonian, Finnish, French, German, Greek, Hungarian, Irish, Italian, Latvian, Lithuanian, Maltese, Polish, Portuguese, Romanian, Slovak, Slovenian, Spanish, Swedish, Arabic, Catalan, Chinese, Galician, Hindi, Japanese, Korean, Norwegian, Russian, Turkish, and Ukrainian. " "The EuroLLM project has the goal of creating a suite of LLMs capable of understanding and generating text in all European Union languages as well as some additional relevant languages."

Table B.1: Model selection, including direct links to Hugging Face model cards, and direct quotes from each model’s model card related to multilingual capabilities. The 5 models which were selected for evaluation on the test set are highlighted in bold.

## B.2.2 Decoding and constrained generation

Following existing work (Bavaresco et al., 2025; Törnberg, 2024; Pavlovic and Poesio, 2024), we apply greedy decoding across all model-prompt combinations. For reasoning steps (in UP3 and UP4), we initially considered using each model’s default and/or explicitly recommended sampling parameters, but did not observe consistent performance improvement.

To ensure that the LLM produces a valid category in response to classification queries, we use the `outlines` library (Willard and Louf, 2023) wrapping around transformers generation. For UP1, UP1-FS and the second turn in UP3 & UP4, we apply the following output schema:

```
output_schema = Literal[
    "NO CREDIBILITY ASSESSMENT",
    "POSITIVE CREDIBILITY ASSESSMENT",
    "NEGATIVE CREDIBILITY ASSESSMENT"
]
```

For UP2, we apply the following output schema at each turn:

```
# 1st turn:
output_schema = Literal["Y", "N"]

# 2nd turn
output_schema = Literal["POSITIVE", "NEGATIVE"]
```

Five models in our selection were found to have limited support for constrained generation: Mistral, EuroLLM and Bielik models. We therefore apply rule-based logic to extract the chosen category from their outputs (cf. Figure B.4).

```
if set(get_args(output_schema)) == {"Y", "N"}:
    output = output[0]
elif set(get_args(output_schema)) == {"POSITIVE", "NEGATIVE"}:
    if output[0] in ["P", "N"]:
        output = "POSITIVE" if output[0] == "P" else "NEGATIVE"
    else:
        if "POSITIVE" in output and "NEGATIVE" in output:
            if output.count("POSITIVE") > output.count("NEGATIVE"): output="POSITIVE"
            else: output="NEGATIVE"
        else:
            if "POSITIVE" in output: output="POSITIVE"
            else: output="NEGATIVE"
elif set(get_args(output_schema)) == {"POSITIVE CREDIBILITY ASSESSMENT", "NEGATIVE CREDIBILITY ASSESSMENT", "NO CREDIBILITY ASSESSMENT"}:
    if not ("POSITIVE" in output or "NEGATIVE" in output or "NO " in output):
        logging.warning(f"LLM output '{output}' could not be mapped to 3-class prediction.")
        output = "NEGATIVE CREDIBILITY ASSESSMENT"
    else:
        output = "POSITIVE CREDIBILITY ASSESSMENT" if "POSITIVE" in output else "NEGATIVE CREDIBILITY ASSESSMENT" if "NEGATIVE" in output else "NO CREDIBILITY ASSESSMENT"
```

Figure B.4: How we extract the predicted class from the LLM’s response.

For reasoning steps (in UP3 and UP4), we do not constrain the content or length of the output. We set `max_new_tokens` to an arbitrarily large number (100,000).

## B.2.3 Compute infrastructure

We perform inference on local hardware (GeForce RTX 3080 with 10GB of VRAM) as well as remote compute nodes (single A40 with 48GB of VRAM, and single H100 with 80GB of VRAM), depending on the model size. Inference time per sample varied widely depending on the model and prompt combination, ranging from 0.1s to 2 minutes per sample.

## C Validation set evaluation

Figure C.5 shows the performance of individual LLM annotators on the validation set. Each datapoint corresponds to a single model-UP-SP combination, with 630 datapoints in total.



Figure C.5: Classification performance of each model-prompt combination on the validation set in terms of Macro F1 (wrt. the label agreed upon by H1 and H2). Plots are split by user prompt, and color coded by system prompt.

The top 15 LLM annotators are selected by averaging performance (macro F1) for the top-3 UP-SP combinations for each model, and taking the top 5 models  $\times$  3 prompts. Table C.2 shows the performance of the resulting selected LLM annotators on the validation set.

Model	System Prompt	User Prompt	Macro-F1 Score (%)	Cohen’s Kappa	Accuracy (%)
phi-4	SP4	UP2	90.51	0.86	91.43
phi-4	SP3	UP2	89.99	0.86	91.43
phi-4	SP4	UP4	86.34	0.81	88.57
gemma-3-27b-it	SP4	UP1	87.50	0.83	90.00
gemma-3-27b-it	SP3	UP1	84.66	0.80	88.57
gemma-3-27b-it	SP4	UP1-FS	79.91	0.69	82.86
Ministral-3-14B-Instruct-2512	SP4	UP3	86.59	0.78	87.14
Ministral-3-14B-Instruct-2512	SP3	UP2	85.95	0.79	87.14
Ministral-3-14B-Instruct-2512	SP5	UP2	82.69	0.76	85.71
Mistral-Small-3.2-24B-Instruct-2506	SP3	UP1-FS	85.63	0.79	87.14
Mistral-Small-3.2-24B-Instruct-2506	SP5	UP2	84.18	0.79	87.14
Mistral-Small-3.2-24B-Instruct-2506	SP4	UP1	81.95	0.74	84.29
Qwen3-30B-A3B-Instruct-2507	SP2	UP1-FS	85.14	0.79	87.14
Qwen3-30B-A3B-Instruct-2507	SP5	UP2	83.61	0.78	87.14
Qwen3-30B-A3B-Instruct-2507	SP5	UP3	82.48	0.78	87.14

Table C.2: **Validation set** performance of the 15 model-prompt combinations that we select for evaluation on the test set. The performance metrics are with respect to human annotations.

## D Test set evaluation

### D.0.1 Top selected model-prompt combinations

Model	System Prompt	User Prompt	Macro-F1 Score (%)	Cohen’s Kappa	Accuracy (%)
phi-4	SP4	UP2	93.66	0.91	94.00
phi-4	SP3	UP2	93.20	0.89	93.00
phi-4	SP4	UP4	94.69	0.91	94.50
gemma-3-27b-it	SP4	UP1	89.89	0.86	91.50
gemma-3-27b-it	SP3	UP1	86.84	0.83	89.50
gemma-3-27b-it	SP4	UP1-FS	84.39	0.78	86.50
Ministral-3-14B-Instruct-2512	SP4	UP3	91.63	0.87	92.00
Ministral-3-14B-Instruct-2512	SP3	UP2	90.62	0.86	91.00
Ministral-3-14B-Instruct-2512	SP5	UP2	89.05	0.84	90.00
Mistral-Small-3.2-24B-Instruct-2506	SP3	UP1-FS	92.52	0.88	92.50
Mistral-Small-3.2-24B-Instruct-2506	SP5	UP2	90.85	0.86	91.50
Mistral-Small-3.2-24B-Instruct-2506	SP4	UP1	92.30	0.88	92.50
Qwen3-30B-A3B-Instruct-2507	SP2	UP1-FS	88.70	0.85	90.50
Qwen3-30B-A3B-Instruct-2507	SP5	UP2	87.95	0.82	89.00
Qwen3-30B-A3B-Instruct-2507	SP5	UP3	88.70	0.83	89.50

Table D.3: **Test set performance** of the 15 model-prompt combinations that we select for evaluation on the test set. The performance metrics are with respect to human annotations, taking the majority vote between annotators.

## D.0.2 Inter-LLM agreement

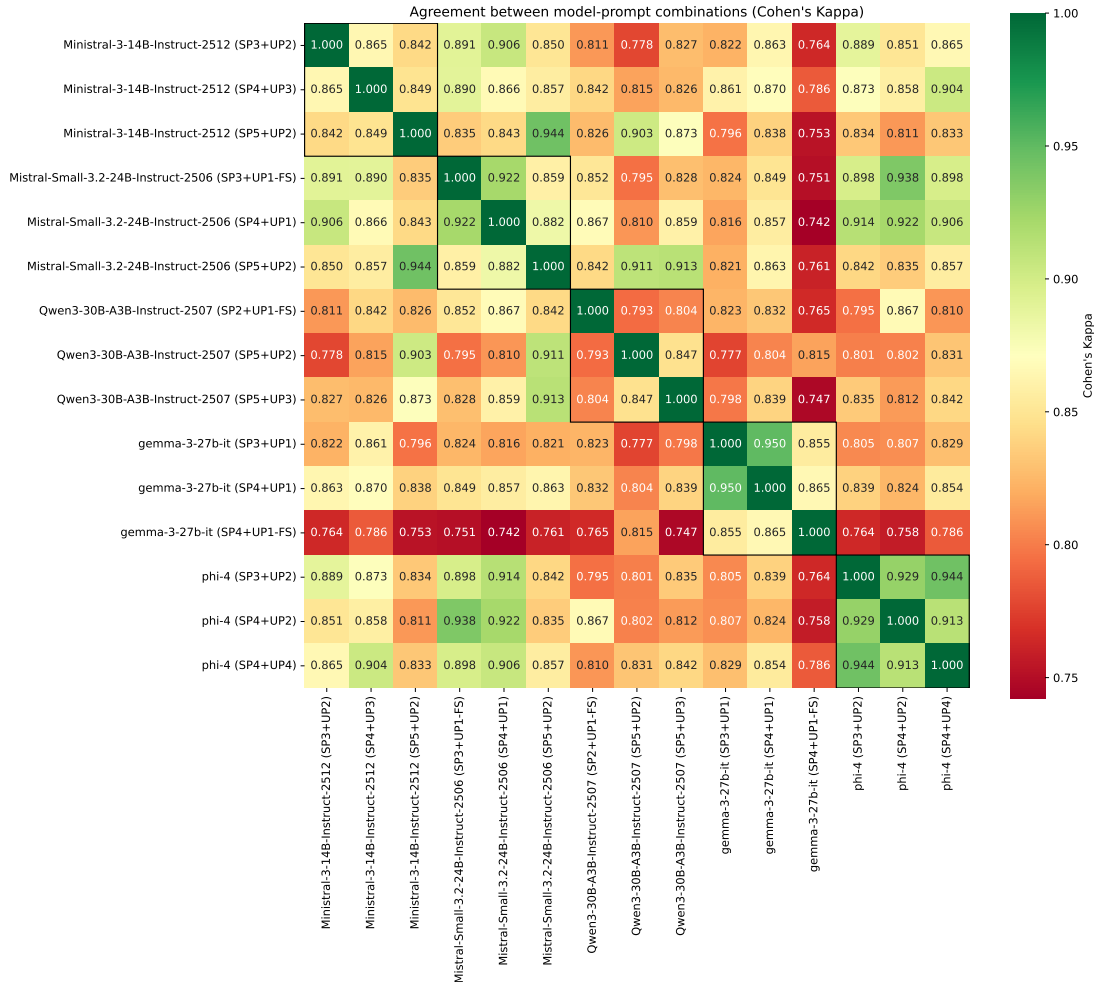


Figure D.6: Cohen's Kappa between pairs of LLM annotators on the test set. Pairs with the same model are outlined in black.

## D.0.3 Sensitivity to prompt and model choice

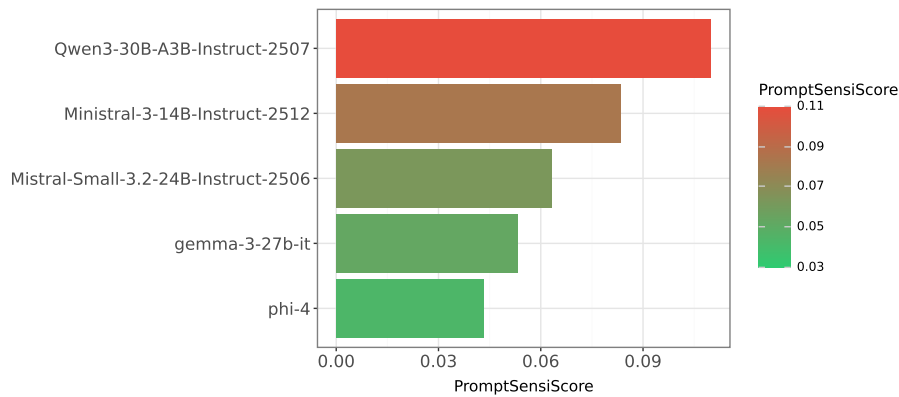


Figure D.7: Instance-level prompt sensitivity across models. Prompt Sensitivity Scores (PSS) across 15 configurations (5 models  $\times$  3 prompt combinations each) on 200 test cases. Phi-4 shows lowest sensitivity (PSS=0.043), gemma-3-27b (PSS=0.053), Mistral-Small-24B (PSS=0.063), Ministral-14B (PSS=0.087), and Qwen3-30B highest (PSS=0.110). Green indicates robust, red indicates sensitive.

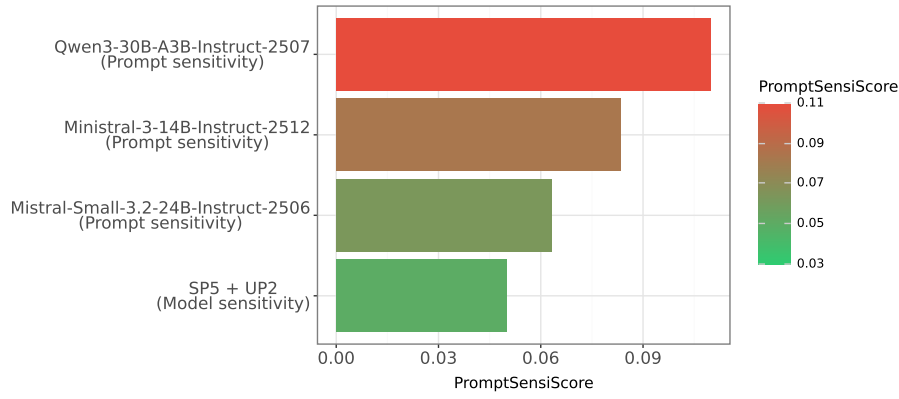


Figure D.8: Model sensitivity versus prompt sensitivity for SP5+UP2. Model sensitivity (green, PSS=0.110) computed across three architectures (Ministral-14B, Mistral-Small-24B, Qwen3-30B) using identical SP5+UP2 prompts on 200 cases, versus prompt sensitivity within each model across three prompts (red/orange/green bars: PSS=0.110, 0.087, 0.063 respectively).

### D.0.4 Inter-class confusion

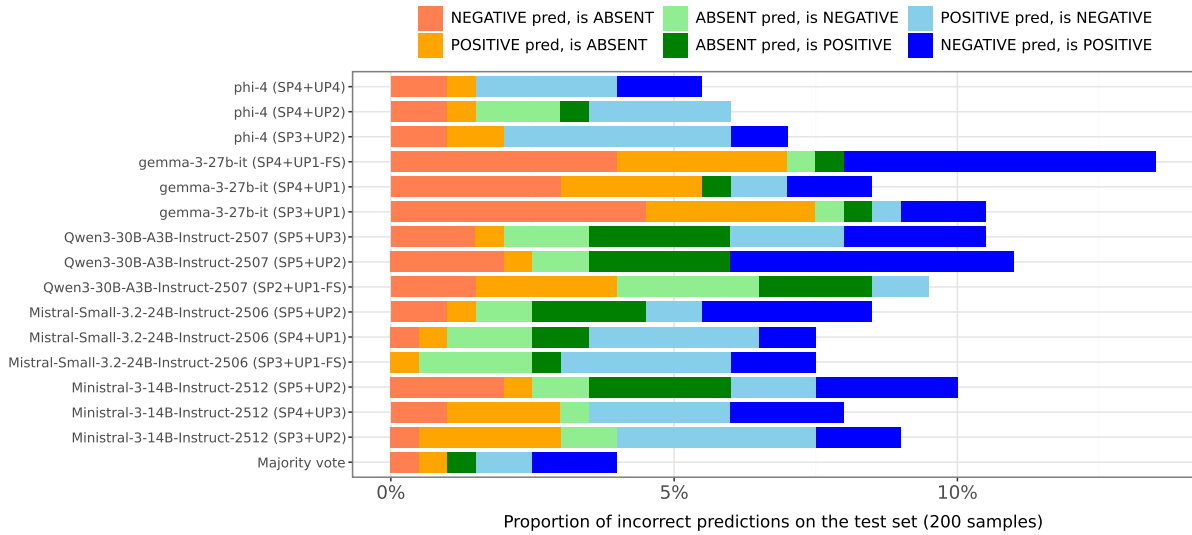


Figure D.9: (Larger version of Figure 9) Mistakes made by individual LLMs and the ensemble (bottom), color-coded by class confusion.

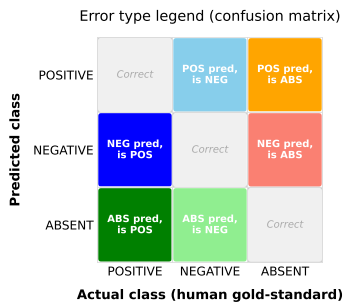


Figure D.10: Legend for Figure D.9

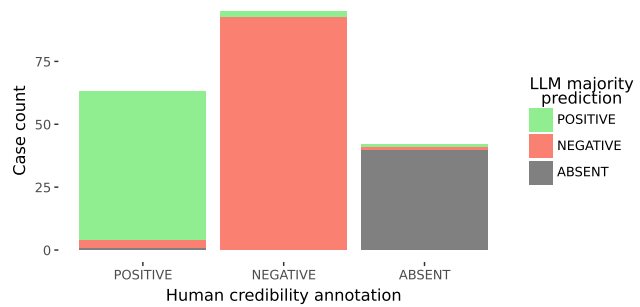


Figure D.11: LLM majority prediction (majority vote across ensemble of 15 LLMs) vs. gold standard human annotation.

## D.0.5 Consistently misclassified cases

H1 was asked to rate mistake severity according to the following 3 categories:

- (A) *I consider this mistake to be an acceptable answer, as I hesitated myself between the LLM's prediction & what I picked when annotating / this mistake is making me rethink/reconsider my own annotation*
- (B) *mistake is understandable and not severe, but I would not expect a fellow domain expert who fully understands the codebook to make it*
- (C) *mistake is severe / not acceptable / calls for an improvement of the model/prompt/etc.*

The two cases (ID 919 and 4317) which are most frequently misclassified across the LLM annotators were rated as (A) by H1, and are shown in Table D.4, along with a reflection by H1 and an LLM's reasoning.

### Case ID 919 (test set)

In January 2023, the Board granted a residence permit (Refugee Convention Status) to a male citizen of Afghanistan. Earlier in 2018, the Board had upheld the Danish Immigration Service's decision to refuse asylum, after which the Board reopened the case in 2022. The person in question had entered the country in 2015. The Refugee Appeals Board stated: "The applicant is an ethnic Hazara and a Shiite Muslim from [city A], [province], Afghanistan. The applicant was born and raised in [city B], Iran. The applicant has not been a member of any political or religious associations or organizations, nor has he been politically active in any other way. The applicant originally cited as his reason for seeking asylum that he feared being killed by the Taliban if he returned to Afghanistan. His family left Afghanistan in 1996 after a conflict with the Taliban, who had killed [several of the applicant's father's family members]. During the reopening of the case, the applicant cited as his reason for seeking asylum that, as a result of the change of power in Afghanistan and the current situation in the country, he would be at risk of persecution because he would be considered Westernized. The applicant further stated that he is of Hazara ethnicity and a Shiite Muslim.

With regard to the applicant's original reason for asylum, which is linked to a derived risk to the applicant due to his family's conflict with the Taliban, reference is made to the Refugee Board's decision of [autumn] 2018 and the reasons stated therein. The fact that the Taliban has now taken power in Afghanistan cannot lead to a different assessment. It follows from the above that this reason for asylum cannot form the basis for asylum under Section 7 of the Aliens Act. The conditions for Westernized persons and persons of Hazara ethnicity (and Shiite Muslims) and the risk assessment to be made in this connection in the event of a return to Afghanistan are described in the background material, inter alia, in sections 3.13 and 3.14 of EUAA, Country Guidance: Afghanistan, January 2023. The Refugee Appeals Board notes in this connection that the decisive factor in the asylum assessment in cases involving questions of so-called Westernization is whether, based on an overall assessment, the applicant has such experience of life in Afghanistan that he or she will be able to resume their life in the country in such a way that he or she will not attract the attention of the Taliban. In this regard, it should be noted that the applicant entered Denmark as an unaccompanied minor in [autumn] 2015, when he was [12-15 years old], and that he has therefore spent a significant part of his schooling in [middle school and high school] and his youth in Denmark, where he has lived as a Danish teenager, including the personal challenges he has explained in the lawyer's statement and during the board meeting. The applicant speaks fluent Danish, and the board meeting was conducted in Danish. Based on the applicant's appearance at the board meeting and the information provided, it must be assumed that the applicant has adopted a Western lifestyle. Since the applicant, according to his explanation, which the Board has taken as a basis, has never lived in Afghanistan, was born and raised in Iran, speaks Dari with an Iranian dialect, and has neither family nor other network in Afghanistan that could support and help him upon return, the Refugee Board finds – also taking into account that the applicant is of Hazara ethnicity (and a Shiite Muslim) – that, based on an overall assessment, it must be considered probable that the applicant will not be able to establish himself in Afghanistan in such a way that he does not attract the attention of the Taliban. Furthermore, and since the applicant is of Hazara ethnicity, as stated, the Refugee Appeals Board finds, after assessing the overall circumstances of the case, that the applicant has demonstrated that he would be at risk of persecution covered by section 7(1) of the Aliens Act if he returned to Afghanistan. The Refugee Board therefore grants the applicant a residence permit pursuant to section 7(1) of the Aliens Act. Reference number: Afgh/2023/8/MKTO.

Original case text: <https://fln.dk/praksis/2023/februar/afgh20238/>

**Annotator H1:** Q1: Credibility assessment present? **No** (Confidence: Low)

**Annotator H2:** Q1: Credibility assessment present? **No** (Confidence: Medium)

**15 LLM annotators:** Positive credibility assessment **X**

**H1's reflection:** I think the issue could stem from the part around here: "*Herefter, efter ansøgerens fremtræden under nævnsmødet og det i øvrigt oplyste må det lægges til grund, at ansøgeren har tilegnet sig en vestlig livsstil. Da ansøgeren samtidig efter sin forklaring, som nævnet har lagt til grund, aldrig har boet i Afghanistan, er født og opvokset i Iran, taler dari med iransk dialekt og hverken har familie eller øvrigt netværk i Afghanistan, der ville kunne støtte og hjælpe ham ved en tilbagevenden, finder Flygtingenævnet - også under hensyntagen til, at ansøgeren er hazara af etnicitet (og shiamuslim) - at det efter en samlet vurdering må anses for sandsynliggjort, at ansøgeren ikke vil kunne etablere sig i Afghanistan på en sådan måde, at han ikke påkalder sig Talibans opmærksomhed.*"

First of all, there is the part about "lægges til grund" (usually used a lot for the credibility) that the claimant has acclaimed a Western lifestyle. I think for me the core of this case was the probability of the claimant to establish a life without attracting the attention of the Taliban and whether the person had acquired a Western enough lifestyle. The style of writing for me generally came down to both a future oriented judgment of what is Western enough or different enough to reach a threshold

of potentially attracting attention and hence being at risk of persecution (more the direction of risk assessment), and less on whether the identity of the claimant can be accepted or past events described are accepted as the basis of the testimony, and I think therefore the core of the issue discussed in this case went for me into a different kind of assessment than classic, text-book credibility of "do we *lægge til grund* what you are saying about yourself and your past experiences to use this as a basis to continue further assessment on". I can imagine that an LLM does not distinguish that necessarily as a categorically different assessment, and it might be very well that I went astray here on the human side.

Further, there is this second part that I copied in, about the accepting that the claimant never lived in Afghanistan, grew up in Iran etc. and reading this part I do think very much that this alone I would label as a credibility assessment, and I think that this is why I scored with low confidence on these cases, and absolutely can see how I as a human took one too many assumptions."

---

#### **Reasoning output of phi-4, prompted with SP4+UP4:**

The task is to determine whether a credibility assessment is made in the decision of the Danish Refugee Appeals Board (Flygtningenævnet) and, if so, to classify its sentiment as either positive, negative, or no credibility assessment. The decision involves a male national from Afghanistan who was granted a residence permit (K-status) in January 2023. The case involves assessing the risk of persecution upon return to Afghanistan, considering the applicant's background as an ethnic Hazara Shia Muslim, his upbringing in Iran, and his life in Denmark.

#### **### Step 2: Preliminary Judgment**

Upon reviewing the case text, it appears that the decision does involve a credibility assessment. The decision explicitly states that the applicant's explanation about not having lived in Afghanistan, being raised in Iran, and having a Western lifestyle is accepted ("*nævnet har lagt til grund*"). This acceptance of the applicant's account suggests a positive credibility assessment, as the core of his narrative is believed, allowing the risk assessment to proceed with his information.

#### **### Step 3: Critical Assessment**

To ensure accuracy, let's critically assess the decision:

- **Credibility Assessment Indicators:** The decision uses phrases like "*nævnet har lagt til grund*," indicating that the applicant's explanation is accepted as true. This aligns with the criteria for a positive credibility assessment.
- **Risk Assessment:** The decision proceeds to assess the risk of persecution based on the accepted account, considering the applicant's Western lifestyle and lack of ties in Afghanistan.
- **No Explicit Rejection:** There is no language indicating that the core of the applicant's account is rejected or deemed non-credible.

Given these points, the preliminary judgment of a positive credibility assessment seems consistent with the decision's content.

#### **### Step 4: Confirm Final Answer**

The final classification is **POSITIVE CREDIBILITY ASSESSMENT**. The decision explicitly accepts the core of the applicant's account, allowing the risk assessment to proceed with his information. The decision does not indicate any rejection of the core narrative, and the applicant's explanation is used as a basis for further analysis.

#### **### Step 5: Evaluate Confidence**

I am confident (90%) in this analysis. The decision clearly uses language that indicates acceptance of the applicant's core narrative, which is a key indicator of a positive credibility assessment. The decision's structure and content align well with the criteria for this classification. The remaining uncertainty (10%) stems from the complexity of legal language and the potential for nuanced interpretations, but the evidence strongly supports the conclusion.

### **Case ID 4317 (test set)**

In November 2018, the Board upheld the Danish Immigration Service's decision regarding a male national from an unknown country of origin. He entered the country in 2015. The Refugee Board stated: "The applicant is an ethnic Arab and a Shiite Muslim from Al Jahra, Kuwait. The applicant has not been a member of any political or religious associations or organizations, nor has he been politically active in any other capacity. As grounds for asylum, the applicant has stated that he fears being imprisoned by the Kuwaiti authorities upon return to his home country due to his illegal departure. The applicant also fears the general conditions for Bidoons. In support of his asylum claim, the applicant has stated that he is an unregistered Bidoon and therefore has no right to a birth certificate, an ID card, schooling, or health care in Kuwait. The applicant has further stated that his uncle, [A], participated in a demonstration [in early] 2014. [In early] 2014, [A] was arrested at the family's shared residence and was subsequently imprisoned until [spring] 2014. After Bader's departure, the Kuwaiti authorities visited the applicant's residence an unknown number of times, causing the applicant's family to feel harassed. Against this background, the applicant left Kuwait in mid-[fall] 2015 together with his family. In the mother's case, the Refugee Board has reached the following decision: "The Refugee Board finds that it cannot rely on the applicant's statement regarding her identity, her nationality, or her grounds for asylum. In this connection, it is noted that the applicant has demonstrated limited knowledge of her home area, where she claims to have lived her entire life. For instance, she has been unable to identify street addresses in her neighborhood other than her own, nor is she familiar with the location of her father's grave in relation to her residence. The documents submitted by the applicant do not lead to a different assessment of the credibility of her statement. Furthermore, she has provided conflicting accounts regarding visits to the authorities following her brother's release, whose asylum claim was, incidentally, rejected by the Refugee Board in the fall of 2018. Against this background, the Board does not find that the applicant has established a likelihood that, upon return to her home country, she would be at a concrete or individual risk of persecution or abuse covered by Section 7 of the Aliens Act. It is added that the applicant has not reported having had conflicts in other countries, including not in Iraq. It is noted that the general conditions for Bidoon in Kuwait, viewed in isolation, do not constitute grounds for asylum, and that the

sanctions for illegal departure, based on the available background information, are not considered disproportionate. The Refugee Board therefore upholds the Danish Immigration Service's decision." Since the applicant's grounds for asylum correspond to those stated by his mother, and since the circumstances included in the decision regarding the applicant's identity, including nationality, are identical to those of his mother, the Refugee Board, with reference to this decision, also finds that this applicant has not established a likelihood that, upon return to his home country, he would be at a concrete and individual risk of persecution or abuse covered by Section 7 of the Aliens Act. The Refugee Board therefore upholds the Danish Immigration Service's decision." Hjem/2018/88/CHA

Original case text: <https://fln.dk/praksis/2019/april/hjem201888/>

**Annotator H1:** Q1: Credibility assessment present? **No** (Confidence: Low)

**Annotator H2:** Q1: Credibility assessment present? **No** (Confidence: Low)

**14 LLM annotators:** Negative credibility assessment ✗

**1 LLM annotator:** No credibility assessment ✓

**H1's reflection for case 4317:** This might also be confusing to other human annotators and I am actually inclined to think that I should have labeled this differently (hence low confidence): The reasoning expands upon the credibility of the mother's explanation, and rejects the mother's testimony due to credibility issues. We had originally thought to exclude references to credibility assessments of family members, in cases where the claimant in the case to be annotated was only rejected or accepted based on the credibility of the family member (as this entails no credibility assessment explicitly on the claimant). However, in this particular case, I would say that they almost reject the claimant because they have the same motive as the mother, who is seen as not credible in her motive, i.e. I actually think one could argue that the claimant themselves is also perceived as non-credible. And from an LLM perspective I think there is so much language about credibility, that I am super interested in which combination actually ruled against this being credibility.

**Reasoning output of phi-4, prompted with SP4+UP4:**

**### Step 1: Clarify Understanding of the Task and the Case**

The task is to analyze a decision from the Danish Refugee Appeals Board (Flygtningenævnet) to determine if a credibility assessment was made and, if so, its sentiment. The decision involves a male applicant from Kuwait, who claims asylum based on fears of persecution due to his status as an "unregistered" Bidoon and his family's experiences with Kuwaiti authorities. The decision must be classified into one of three categories: NO CREDIBILITY ASSESSMENT, POSITIVE CREDIBILITY ASSESSMENT, or NEGATIVE CREDIBILITY ASSESSMENT.

**### Step 2: Preliminary Judgment**

Upon reviewing the case text, it is evident that the Flygtningenævnet explicitly states that they "find not to be able to base the applicant's explanation regarding his identity, nationality, or asylum motive." This indicates a direct assessment of the applicant's credibility. The decision mentions inconsistencies and a lack of detailed knowledge about his home area, which further supports the conclusion that the core of the applicant's account is not accepted. Therefore, the preliminary judgment is that there is a NEGATIVE CREDIBILITY ASSESSMENT.

**### Step 3: Critical Assessment**

To ensure accuracy, let's critically assess the decision:

- **Explicit Statements:** The decision explicitly states that the applicant's explanations about identity, nationality, and asylum motive cannot be accepted. This is a clear indication of a credibility assessment.
- **Inconsistencies and Lack of Knowledge:** The decision highlights the applicant's limited knowledge of his home area and inconsistencies in his account, which are reasons for rejecting the core of his narrative.
- **Peripheral Facts:** While some peripheral facts might be accepted, the core elements necessary for asylum are not believed, aligning with a NEGATIVE CREDIBILITY ASSESSMENT.

**### Step 4: Confirm Final Answer**

The final answer is NEGATIVE CREDIBILITY ASSESSMENT. The decision explicitly rejects the core of the applicant's account, citing inconsistencies and a lack of detailed knowledge, which are critical for establishing credibility.

**### Step 5: Evaluate Confidence**

I am 95% confident in this analysis. The decision clearly articulates reasons for rejecting the applicant's core narrative, which aligns with the criteria for a **NEGATIVE CREDIBILITY ASSESSMENT**. The explicit language used in the decision leaves little room for alternative interpretations regarding the credibility assessment. The remaining 5% uncertainty accounts for any potential nuances in the interpretation of the decision's language that might not be immediately apparent.

Table D.4: Translation of two case texts from the test set (ID 919 and 4317) which were misclassified by all or almost all 15 LLM annotators, and the corresponding annotation independently assigned by the two domain experts. Under each case text, we include the reasoning output of the best-performing LLM annotator on the test set (bold added for clarity). Case text translated with DeepL.com (free version).

# Cracks in the Bridge—or A Bridge Too Far? Comparing Human and LLM Errors in the Annotation of Bridging Anaphora

Lauren Levine and Amir Zeldes

Georgetown University  
Department of Linguistics  
{le176, amir.zeldes}@georgetown.edu

## Abstract

In this paper, we perform an error analysis on human and LLM annotation data from the recent GUMBridge corpus for varieties of bridging anaphora. We explore the distribution of precision and recall errors made by annotators and how that distribution correlates with bridging subtypes. We find that while LLMs perform substantially worse than human annotators, they are more balanced in their precision and recall scores than humans, whose performance strongly favors precision. With regard to subtypes, we find that *COMPARISON* and *MERONOMY* relations are easier to reliably annotate than the more broadly construed *ENTITY* relations for both human and LLM annotators, but that LLM errors are more distributed across subtypes than human errors. Analyzing these results, we provide insights for future annotation projects on bridging anaphora.

## 1 Introduction

A “bridging anaphor” is a newly introduced entity whose referent is inferable specifically due to its relation to a previously introduced entity in the discourse. Consider the following example:

- (1) There is a house. **The door** is red.<sup>1</sup>

In the example above, the entity “**the door**” (the bridging anaphor) is specifically understood to be the door of the aforementioned “house” (the associative antecedent).

Bridging has been shown to be a difficult phenomenon to annotate, as it is dependent on the annotator’s subjective understanding of entities and entity relations in a discourse (Levine and Zeldes, 2025). Analyzing the different errors made by human and LLM annotators allows us to gain insights into the strengths and limitations of each annotation approach, which can be leveraged in future

efforts for the annotation of bridging anaphora, including the formulation of prompts for LLM-based approaches.

In this short paper, we investigate the following questions to better understand the errors that occur when annotating bridging:

- RQ1** What types of errors (precision vs recall) dominate human and LLM annotation of bridging anaphora?  
**RQ2** How do sub-varieties of bridging anaphora correlate with precision and recall errors in human and LLM annotations?

## 2 Related Work

“Bridging” occurs when a discourse participant constructs an implicature from the entity they are currently processing back to an antecedent entity (Clark, 1975). Bridging has been studied from a variety of theoretical perspectives (e.g., Hawkins, 1978; Prince, 1981; Asher and Lascarides, 1998; Baumann and Riester, 2012), and linguistic resources have been constructed for various languages, including English (Markert et al., 2012, Rösiger, 2018, Poesio and Artstein, 2008; Uryupina et al., 2019), German (Schweitzer et al., 2018, Eckart et al., 2012), Polish (Ogrodniczuk and Zawistawska, 2016), and Czech (Nedoluzhko et al., 2009). While there have been various neural and rule-based systems for the identification of bridging instances (Rösiger et al., 2018, Yu and Poesio, 2020, Kobayashi et al., 2022), and other studies of LLMs’ discourse capabilities (Li and Carenini, 2026), studies of LLMs’ abilities to detect bridging anaphora are limited, with minimal benchmarks available (Bu et al., 2025).

## 3 Data

For our investigation of errors in the annotation of bridging anaphora, we leverage data from the

<sup>1</sup>Bridging anaphora are marked in bold face, and their associative antecedents are underlined.

GUMBridge corpus, a recent effort to annotate sub-varieties of bridging anaphora in English (Levine and Zeldes, 2026a). Built on top of the multi-genre GUM corpus (Zeldes, 2017), which includes pre-existing entity mention and coreference annotations (see Zeldes 2022 for detailed discussion), it uses an information status based definition of bridging, and a schema of sub-varieties with 3 main categories for bridging relations: COMPARISON relations, ENTITY relations, and SET relations. The size of the partitions of GUMBridge (in terms of tokens and gold bridging instances) are shown in Table 1 (see Levine and Zeldes 2026b for a descriptive analysis of the corpus).

In order to investigate human errors in annotation, we compare data from Levine and Zeldes 2026a’s existing inter-annotator agreement (IAA) study on the annotation of the GUMBridge dev set with the final adjudicated gold version of the dev set. In the IAA study, the 32 documents in the dev set (~30k tokens) were double annotated by pairs of annotators. These annotators were provided with extensive annotation guidelines and participated in several hours of training prior to completing the annotation task. Please see Levine and Zeldes 2026a for further details on the annotation procedure.

To investigate LLM errors in annotation, we compare LLM predictions on the GUMBridge test set with the corresponding gold annotations. Levine and Zeldes 2026a provides several LLM baselines for bridging resolution on the test set, providing scores for several models. We use the judgments from the best performing model, GPT-5, for our error analysis. We also provide results for another SoTA LLM, gemini-3.1-pro-preview, using the same prompts/workflow as in Levine and Zeldes 2026a, in order to have double annotated SoTA LLM data to analyze. A brief summary of this pipeline and the prompts used is included in Appendix C.

All data analyzed in this paper is publicly available with the GUMBridge v1 release, including our additional LLM annotations.<sup>2</sup>

#### 4 RQ 1: Distribution of Error Types for Human and LLM Annotators

Table 2 shows the performance of human and LLM annotators on the identification of bridging instances, i.e., the identifying the bridging anaphora

<sup>2</sup>GUMBridge data, code, and LLM results publicly available: <https://github.com/lauren-lizzy-levine/gumbridge/>

GUMBridge v1	Tokens	Bridging Instances	Bridging per 1k Tokens
Train	213k	4k	18.9
Dev	30k	732	24.5
Test	30k	562	18.6
Test2	18k	379	21.2
Total	291k	5.7k	19.6

Table 1: Distribution of bridging instances across GUMBridge partitions.

	Precision	Recall	F1-Score
<b>Humans (dev)</b>			
Annotator 1	68.5	56.4	61.9
Annotator 2	67.2	44.8	53.8
Avg.	67.9	50.6	57.9
<b>LLMs (test)</b>			
GPT-5	23.2	20.3	21.7
Gemini	<b>27.6</b>	<b>28.9</b>	<b>28.2</b>
Avg.	25.4	24.6	25.0

Table 2: Performance of human and LLM annotators identifying bridging pairs (on dev and test respectively).

**and** correctly resolving it back to its associative antecedent. Partial matches, i.e., correctly identifying a bridging anaphor but not identifying the correct associative antecedent, are counted as incorrect. Average and individual scores from the double annotation of the dev set are reported, as are the average and individual performances of the 2 LLMs on the test set. As one would expect, we observe that the performance of LLMs is substantially weaker than human annotators (avg. F1-score  $\Delta$  of 32.9). This aligns with recent work showing that LLMs excel at tasks which are solvable by less trained crowd workers, but often fail to reach the level of skilled annotators on complex tasks (Kasner et al., 2026). Additionally, we observe that the precision and recall scores pattern differently for humans and LLMs. Humans have notably higher precision scores than recall scores (avg.  $\Delta$  of 17.3), while the precision and recall scores of the LLM baselines are much closer (avg.  $\Delta$  of 0.8).

To further explore the distribution of error types, in Figure 1 we show the distribution of False Positive (FP) precision errors and False Negative (FN) recall errors for humans and LLMs, including whether one or both annotators made the same er-

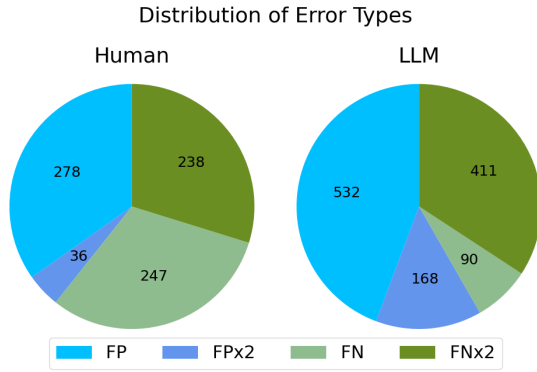


Figure 1: Distribution of error types for human and LLM annotators. FP and FN refer to errors committed by a single annotator, while FPx2 and FNx2 refer to errors committed by both annotators.

ror.<sup>3</sup> The FP and FN categories are for instances where an error was committed by only a single annotator, while the FPx2 and FNx2 categories are for instances where the same error was committed by both annotators. We once again see that human annotators have a higher proportion of recall errors, while the LLMs now have a higher proportion of precision errors. This indicates that the two LLMs were more varied in their False Positives than their False Negatives, and that they are more likely to over-generate than their human counterparts.

Additionally, we note that the human False Negatives are roughly evenly split between being missed by one annotator or both annotators, while the LLM False Negatives are dominated by instances missed by both models. This suggests that having multiple human annotators will be more valuable than having multiple LLM annotators when trying to have broad coverage for identifying bridging anaphora, as one human annotator is more likely than an LLM to notice what another annotator has missed. However, variation in LLM errors suggests that querying a broader range of models remains valuable, as this diversity is essential for ensembling/voting approaches to bridging anaphora identification.

## 5 RQ 2: Analysis of Subtype Errors for Human and LLM Annotators

Every instance of bridging in the GUMBridge corpus has a subtype annotation which contains one or more subtype labels. These subtype labels fall into three main categories:

<sup>3</sup> $\chi^2$  for error type (FP, FPx2, FN, FNx2) and annotator type (human, LLM) is significant: X-squared = 212.09, df = 3, p-value < 2.2e-16

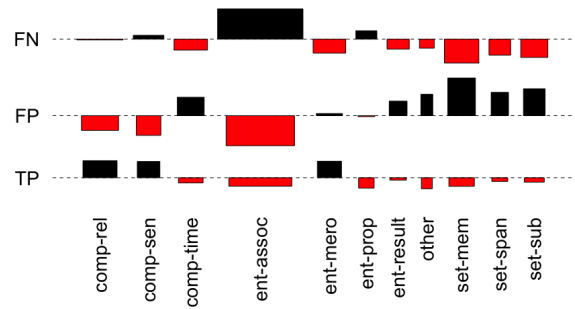


Figure 2:  $\chi^2$  residuals for human error type (False Negatives (FN), False Positives (FP), True Positives (TP)) and bridging subtype label. (X-squared = 168.73, df = 20, p-value < 2.2e-16)

**COMPARISON Relations** The anaphor is preceded by a descriptor which implies a comparison to the antecedent (or vice versa).

- (2) I think her dog is nice, but I want to get a **bigger dog**.

**ENTITY Relations** The anaphor is an attribute or associated entity of the antecedent (or vice versa).

- (3) There is a library around the corner. **The books** are fantastic.

**SET Relations** There is a set/subset relation between the bridging anaphor and antecedent.

- (4) My niece got several toys for her birthday. Her favorite is **the doll**.

Within these three main categories, GUMBridge distinguishes 10 sub-varieties, and there is an additional OTHER category for a total of 11 sub-varieties (see Appendix A for details).

In this section, we investigate whether the subtype(s) of a bridging instance influences whether a human or LLM annotator will identify it. To do this, we examine the distribution of bridging subtype labels and observed error types (including true positives) for human and LLM annotators. Subtype labels for True Positives and False Negatives are taken from the gold labels in GUMBridge dev and test. The False Positive labels are taken from annotator judgments.<sup>4</sup> In Figure 2, we show the residuals from a  $\chi^2$  test for human error type and

<sup>4</sup>False Positives are not included in the LLM analysis because the LLM pipeline was divided into subtasks, and the LLMs were not queried on subtype classification for FPs.

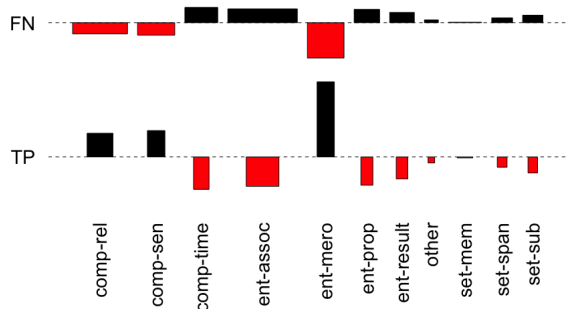


Figure 3:  $\chi^2$  residuals for LLM error type (False Negatives (FN), True Positives (TP)) and bridging subtype label. (X-squared = 87.16, df = 10, p-value = 1.955e-14)

bridging subtype label, and in Figure 3, we show the same for LLM error type.<sup>5</sup>

Looking at Figures 2 and 3, we see that both human and LLM annotators have a higher proportion of COMPARISON-RELATIVE, COMPARISON-SENSE, and ENTITY-MERONOMY in their True Positives, which indicates that these subcategories are more reliably identified by both humans and LLMs. This tendency likely reflects the fact that these subtypes frequently have overt markers which make them easier to recognize. For instance, consider the following example of COMPARISON-RELATIVE/COMPARISON-SENSE:

- (5) I just had a piece of cake, and I want to have **another one**.

In the example above, “another” is a comparative marker which helps to identify the example as an instance of COMPARISON-RELATIVE, and “one” is a common lemma in instances of sense anaphora, which helps to identify the example as an instance of COMPARISON-SENSE. Both items appear in the guidelines and in LLM prompt examples. Looking at the False Negatives row, we also see that instances labeled ENTITY-ASSOCIATIVE are the more common ones to overlook for both humans and LLMs. This is unsurprising, as associative entity relations comprise the broadest sub-variety, covering a variety of implicit relations which lack overt markers, such as relational nouns (e.g., a business → **the customer**), implicit arguments (e.g., a murder → **the victim**), and prototypical associations (e.g., a wedding → **the reception**).

When we look at the human False Positives row

<sup>5</sup>The raw counts of the subtype labels that appear in each error category are given in Appendix B

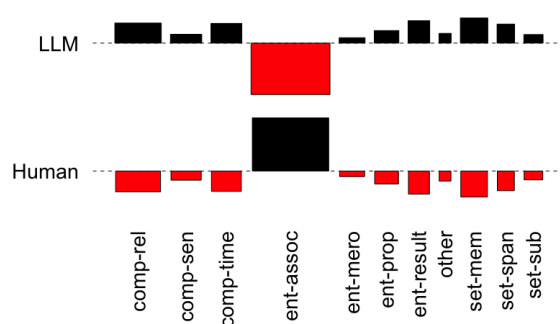


Figure 4:  $\chi^2$  residuals for human and LLM False Negatives and bridging subtype label. (X-squared = 61.472, df = 10, p-value = 1.906e-09)

in Figure 2, we can see that COMPARISON-TIME relations and the SET relations are more commonly mistaken for bridging by human annotators. This is likely because these semantic relations can occur without being anaphoric (i.e., they can be interpreted without reference to the antecedent) but are easily recognizable. As such, they may get flagged by annotators even if they don’t fit the anaphoric criteria for bridging anaphora. Consider the following example:

- (6) The 20th century was a time of rapid advancement. In particular, **the 1960s** were an eventful period.

While the above is semantically an example of a SET-SPAN-INTERVAL/COMPARISON-TIME relation, a discourse participant does not require “the 20th century” to interpret “the 1960s”, which is understandable by itself (though by contrast, “the 60s” could bridge from a particular century).

While the general patterns for which subtype relations are easier are relatively consistent between LLMs and human annotators, we can also compare the two directly by looking at the subtype label distribution of the False Negative occurrences for humans and LLMs. In Figure 4, we show the residuals from a  $\chi^2$  test for human and LLM False Negatives and bridging subtype label. We see very clearly that human errors of omission are concentrated on instances of ENTITY-ASSOCIATIVE, while the errors by LLMs are more spread across subtypes. However, a plurality of missed bridging instances for the LLMs are ENTITY-ASSOCIATIVE (see Appendix B), and as Figure 3 shows ENTITY-ASSOCIATIVE is actually a difficult subtype category for LLMs,

just as it is for humans. As such, it is not the case that LLMs perform very well on associative bridging instances, just that their errors are more spread across the subtypes, while humans are more strongly concentrated on ENTITY-ASSOCIATIVE.

## 6 Conclusion: Takeaways for Future Annotation Efforts

In this paper, we compared human and LLM errors in the annotation of bridging anaphora. Looking at the error distributions of human and LLM annotators, we saw that:

- LLM annotators are worse overall
- Humans favor precision over recall, while LLM are more balanced between the two
- Human variability can provide broader coverage of bridging than LLMs

Looking at the subtype distributions in the errors of human and LLM annotators, we saw that:

- Human and LLM annotators find the same bridging subtypes easier to identify
- Humans are prone to incorrectly identifying SET relations as bridging instances
- Relative to humans, LLM errors are more distributed across bridging subtypes

Based on these findings, we see that for the time being there is a clear advantage to using human annotators when creating data for bridging anaphora, and that human annotation projects should focus on finding a way to boost recall. In this regard, LLMs may be useful for selecting candidates for human annotators to consider as bridging anaphora. LLMs may also be useful in flagging potential human errors due to inattention or misapplication of guidelines (see [Nahum et al. 2025](#); [Chochlakis et al. 2025](#)). Finally, future annotation guideline refinement should focus on further specifying the ENTITY-ASSOCIATIVE subtype, as it is the most difficult sub-variety for both human and LLM annotators, as well as adding attention to non-anaphoric set relations being mistaken for bridging anaphora.

## Limitations

This effort focuses on performing error analysis on annotation data produced by a previous paper. As such, it is constrained by the limitations of the

data produced in that work (such as LLM False Positives not having subtype annotations) and does not create new bridging data besides an additional LLM baseline. However, the comparison of human and LLM errors in the annotation of bridging anaphora is previously unexplored, and we hope the results of this analysis provide novel insights for annotation work.

## References

- Nicholas Asher and Alex Lascarides. 1998. Bridging. *Journal of Semantics*, 15(1):83–113.
- Stefan Baumann and Arndt Riester. 2012. Referential and lexical givenness: Semantic, prosodic and cognitive aspects. *Prosody and meaning*, 25:119–162.
- Lanni Bu, Lauren Levine, and Amir Zeldes. 2025. DiscoTrack: A multilingual LLM benchmark for discourse tracking. *arXiv preprint arXiv:2510.17013*.
- Georgios Chochlakis, Peter Wu, Tikka Arjun Singh Bedi, Marcus Ma, Kristina Lerman, and Shrikanth Narayanan. 2025. [Humans hallucinate too: Language models identify and correct subjective annotation errors with label-in-a-haystack prompts](#). In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 19637–19656, Suzhou, China. Association for Computational Linguistics.
- Herbert H. Clark. 1975. [Bridging](#). In *Theoretical Issues in Natural Language Processing*.
- Kerstin Eckart, Arndt Riester, and Katrin Schweitzer. 2012. [A discourse information radio news database for linguistic analysis](#). In *Linked Data in Linguistics: Representing and Connecting Language Data and Language Metadata*, pages 65–76, Berlin, Heidelberg. Springer Berlin Heidelberg.
- John A. Hawkins. 1978. Definiteness and indefiniteness: A study in reference and grammaticality prediction. *Journal of Linguistics*, 27:405–442.
- Zdeněk Kasner, Vilém Zouhar, Patrícia Schmidtová, Ivan Kartáč, Kristýna Onderková, Ondrej Platek, Dimitra Gkatzia, Saad Mahamood, Ondrej Dusek, and Simone Balloccu. 2026. [LLMs as span annotators: A comparative study of LLMs and humans](#). In *Proceedings of the First Workshop on Multilingual Multicultural Evaluation*, pages 1–22, Rabat, Morocco. Association for Computational Linguistics.
- Hideo Kobayashi, Yufang Hou, and Vincent Ng. 2022. [Constrained multi-task learning for bridging resolution](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 759–770, Dublin, Ireland. Association for Computational Linguistics.

- Lauren Levine and Amir Zeldes. 2025. [Subjectivity in the annotation of bridging anaphora](#). In *Proceedings of the 19th Linguistic Annotation Workshop (LAW-XIX-2025)*, pages 48–59, Vienna, Austria. Association for Computational Linguistics.
- Lauren Levine and Amir Zeldes. 2026a. [Gumbridge: A corpus for varieties of bridging anaphora](#). In *Proceedings of the Fifteenth Language Resources and Evaluation Conference (LREC 2026)*, pages 6823–6837, Palma, Mallorca, Spain. European Language Resources Association (ELRA).
- Lauren Levine and Amir Zeldes. 2026b. What’s in a bridge?: A descriptive, multi-genre analysis of the gumbridge corpus for varieties of bridging anaphora. In *Proceedings of the 2nd Joint Workshop on Computational Approaches to Discourse, Context and Document-Level Inferences (7th CODI) and Computational Models of Reference, Anaphora and Coreference (9th CRAC)*, San Diego, California, USA. Association for Computational Linguistics.
- Chuyuan Li and Giuseppe Carenini. 2026. [BeDiscover: The benchmark of discourse understanding in the era of reasoning language models](#). In *Proceedings of the 19th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4417–4479, Rabat, Morocco. Association for Computational Linguistics.
- Katja Markert, Yufang Hou, and Michael Strube. 2012. [Collective classification for fine-grained information status](#). In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 795–804, Jeju Island, Korea. Association for Computational Linguistics.
- Omer Nahum, Nitay Calderon, Orgad Keller, Idan Szpektor, and Roi Reichart. 2025. [Are LLMs better than reported? detecting label errors and mitigating their effect on model performance](#). In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 26782–26809, Suzhou, China. Association for Computational Linguistics.
- Anna Nedoluzhko, Jiří Mírovský, and Petr Pajas. 2009. [The coding scheme for annotating extended nominal coreference and bridging anaphora in the Prague dependency treebank](#). In *Proceedings of the Third Linguistic Annotation Workshop (LAW III)*, pages 108–111, Suntec, Singapore. Association for Computational Linguistics.
- Maciej Ogrodniczuk and Magdalena Zawistawska. 2016. [Bridging relations in Polish: Adaptation of existing typologies](#). In *Proceedings of the Workshop on Coreference Resolution Beyond OntoNotes (COR-BON 2016)*, pages 16–22, San Diego, California. Association for Computational Linguistics.
- Massimo Poesio and Ron Artstein. 2008. [Anaphoric annotation in the ARRAU corpus](#). In *Proceedings of the Sixth International Conference on Language Resources and Evaluation (LREC’08)*, Marrakech, Morocco. European Language Resources Association (ELRA).
- Ellen F. Prince. 1981. Toward a taxonomy of given-new information. *Radical pragmatics*, pages 223–255.
- Ina Rösiger. 2018. [BASHI: A corpus of Wall Street Journal articles annotated with bridging links](#). In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan. European Language Resources Association (ELRA).
- Ina Rösiger, Arndt Riester, and Jonas Kuhn. 2018. [Bridging resolution: Task definition, corpus resources and rule-based experiments](#). In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 3516–3528, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- Katrin Schweitzer, Kerstin Eckart, Markus Gärtner, Agnieszka Falenska, Arndt Riester, Ina Rösiger, Antje Schweitzer, Sabrina Stehewien, and Jonas Kuhn. 2018. [German radio interviews: The GRAIN release of the SFB732 silver standard collection](#). In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*. European Language Resources Association (ELRA).
- Olga Uryupina, Ron Artstein, Antonella Bristot, Federica Cavicchio, Francesca Delogu, Kepa Joseba Rodríguez, and Massimo Poesio. 2019. [Annotating a broad range of anaphoric phenomena, in a variety of genres: the arrau corpus](#). *Natural Language Engineering*, 26:95 – 128.
- Juntao Yu and Massimo Poesio. 2020. [Multitask learning-based neural bridging reference resolution](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 3534–3546, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Amir Zeldes. 2017. [The GUM corpus: Creating multilayer resources in the classroom](#). *Language Resources and Evaluation*, 51(3):581–612.
- Amir Zeldes. 2022. [Opinion piece: Can we fix the scope for coreference?](#) *Dialogue & Discourse*, 13:41–62.

## A GUMBridge Bridging Subtypes

This appendix briefly details the bridging subtype varieties annotated in the GUMBridge corpus, which are reflected in its guidelines and in the prompts used in the GUMBridge paper to elicit LLM annotations. For the full guidelines, we refer readers to [Levine and Zeldes 2026a](#).

**COMPARISON-RELATIVE** The anaphor is preceded by a comparative marker which implies a comparison to the antecedent (e.g., several women → **other women**).

**COMPARISON-SENSE** The type of the anaphor is omitted but inferable via comparison to the antecedent (e.g., a Chinese restaurant → **the Italian one**).

**COMPARISON-TIME** The anaphor refers to a specific time/time frame which is understandable with reference to the time/time frame expressed by the antecedent (e.g., Wednesday → **yesterday**).

**ENTITY-MERONOMY** The anaphor has a part-whole relation with the antecedent, including physical subparts, substance-portion, and regions/subsections (e.g., a house → **the door**).

**ENTITY-PROPERTY** The anaphor is a physical or intangible property of the antecedent, such as smell, length, or style (e.g., a bouquet of roses → **the scent**).

**ENTITY-RESULTATIVE** The anaphor is logically inferable from the antecedent. This is often the result of a transformative/product producing process, like cooking/baking (e.g., flour → **the bread**).

**ENTITY-ASSOCIATIVE** The anaphor is an attribute or closely associated entity of the antecedent (e.g., a library → **the books**).

**SET-MEMBER** The anaphor is an element of the antecedent set. This includes group-member and class-instance relations (e.g., several books → **the mystery novel**).

**SET-SPAN-INTERVAL** The anaphor is a sub-span of the spatial or temporal antecedent interval (e.g., Sunday → **the morning**).

**SET-SUBSET** The anaphor is a subset of the antecedent set (e.g., a group of students → **the boys**).

**OTHER** The OTHER category is for instances which fit the information status based definition of a bridging pair but do not fall into any of the bridging subtype categories outlined above.

## B Bridging Subtype Label Counts in Error Classes

Table 3 gives the raw counts of the bridging subtype labels that occur in each error class (including True Positives) for the human and LLM annotation of bridging instances. No False Positive counts are provided for the LLM annotations because the LLM pipeline was divided into subtasks (anaphor recognition, antecedent resolution, and subtype

Subtype	Human			LLM	
	TP	FP	FN	TP	FN
COMPARISON-RELATIVE	53	32	70	35	104
COMPARISON-SENSE	26	8	34	19	45
COMPARISON-TIME	17	36	28	1	49
ENTITY-ASSOCIATIVE	118	104	305	25	198
ENTITY-MERONOMY	29	24	24	32	30
ENTITY-PROPERTY	4	9	19	0	30
ENTITY-RESULTATIVE	8	17	11	1	27
SET-MEMBER	16	51	19	9	42
SET-SPAN-INTERVAL	7	20	7	2	18
SET-SUBSET	11	30	12	1	18
OTHER	1	11	4	1	8

Table 3: Human and LLM subtype label counts in the different error classes: True Positives (TP), False Positives (FP), and False Negatives (FN).

classification), and the LLMs were not queried on subtype classification for False Positives.

## C LLM Pipeline and Prompts

In this appendix we provide a brief description of the LLM pipeline used to create the bridging resolution baseline data analyzed in this paper. Please see [Levine and Zeldes 2026a](#) for further details.

The LLMs are queried individually for each of the following bridging resolution subtasks: (1) anaphora recognition, (2) antecedent selection, and (3) subtype categorization. For each subtask, the models are provided with a separate prompt which gives a detailed explanation of the task based on the GUMBridge annotation guidelines, with a series of few-shot examples. For the anaphor recognition subtask, the models are queried sentence by sentence through the document. For the antecedent selection subtask, the models are queried once for each bridging anaphor in the gold annotations. For the subtype categorization subtask, the models are queried once for each bridging pair in the gold annotations. Prompt templates for each subtask are included below.

### C.1 Anaphor Recognition

You are a linguistic analyst whose job is to find cases of bridging anaphora: mentions of newly introduced entities (noun phrases) in a text, for which a reader would need to refer back to a previously mentioned, non-identical entity to resolve their meaning. There are several classes of bridging anaphors, any of which should be identified in the text being analyzed. In the following examples, the bridging anaphor is surrounded by \*asterisks\*.

comparison-relative: The anaphor is preceded by a comparative marker (other, another, same, more, ordinal modifiers, comparative adjectives, superlatives, etc.) which implies a comparison to the antecedent. For example: "The children... \*another child\*" (=another with comparison to the aforementioned children); similar cases may be \*similar children\*, \*older children\* (compared to the aforementioned children), etc.

comparison-sense: the semantic type of a phrase requires a previous mention to identify it, for example "the Italian "restaurant... \*a Chinese one\*" (we can't know "a Chinese one" is a restaurant without referring back to the Italian restaurant), or "\*another one\*", "\*the others\*" etc.

comparison-time: the anaphor refers to a specific time/timeframe which is understandable with reference to the antecedent, for example: "Tuesday, February 2nd ... \*the following week\*"

entity-meronymy: the anaphor is a subunit of the antecedent (part-whole), including physical subunits, portion-substance relations, and regions/subsections. For example: "the house ... \*the door\*" (=of the house).

entity-associative: the anaphor is an attribute or closely associated entity of the antecedent, including both prototypical and inducible associations: "a wedding ... \*the bride\*" (=the bride at that wedding), implicit arguments of a predicate or a verbal nominalization: "a play... \*the performance\*" (=of the play), relational nouns: "a murder ... \*the victim\*"

entity-property: the anaphor is a physical or intangible property of the antecedent (e.g., smell, length, size, style, etc.): "the tea... \*the sweet aroma\*"

entity-resultative: the anaphor is logically inferable from the antecedent (e.g., result, transformation/transmutation, cause): "the dough ... \*the bread\*" (=the dough becomes bread after baking)

set-member: the anaphor is an element of the antecedent set, including groups-member relations and classes-instances: "the cars ... \*the Mazda\*", additionally indefinite members to definite sets: "a candle on each cupcake... \*the candles\*"

set-subset: the anaphor is a subset of the antecedent set: "the cars ... \*the Mazdas\*" (not all Mazdas, just the subset among the aforementioned cars)

set-span-interval: the anaphor is a sub-span of a spatial or temporal interval defined by the antecedent: "last week... \*Wednesday\*" (=Wednesday of last week), "Sunday... \*the morning\*" (=the morning portion of that Sunday)"

other: the anaphor requires a previous entity for interpretation, but it doesn't fit into any of the above categories. This is a rare class.

There are also some exceptions which should NOT be identified as bridging anaphora:

Coreference: If an entity has a previous mention, it cannot be an instance of bridging. For instance, in "Catherine and Henry had their wedding last week. The bride was very beautiful", even though there is an associative relationship between the wedding and the bride, since "the bride" corefers with "Catherine", which has already been introduced to the discourse, "the bride" is not eligible to be an instance of bridging.

Bridging-contained: If the entity one would need to refer back to in order to understand the bridging anaphor is a direct modifier in the noun phrase of the potential bridging anaphor, e.g. "the focus of the story" or "two of them", it should not be annotated as bridging. In other words, the previous antecedent entity must be outside of the nominal phrase containing the anaphor. An entity that is followed by a prepositional phrase or a relative clause is sufficiently qualified and is thus NOT an instance of bridging.

Generics/Situational bridging: Entities that are accessible due to general world knowledge or situational context are not considered instances of bridging, i.e., if it doesn't have a previous associated antecedent entity to be bridging from, it cannot be bridging.

Possession with an explicit possessive: If the potential bridging anaphor contains an explicit possessive which corefers with the associative antecedent entity, no bridging relation is necessary. Explicit coreference between the associative antecedent and the possessive is sufficient (e.g., [Mark]... [his] house → no bridging, coreference between "Mark" and "his"). Contrast this with [the family] ... \*the house\* → bridging, since we cannot interpret which house it is (the house of the family) without referring to "the family", which is outside of the anaphor phrase.

Here are 2 examples of the task:

Please return a list all of the bridging anaphors in the following text in the order in which they appear. Output the anaphor mention phrase exactly as it appears in the text. If there are no bridging anaphors, return an empty list.

Text:

... with their friends to a picnic. The picnic was supposed to take place in

a grove, but the shade wasn't enough, so they had to find a different place. Conny started to say ...

Answer(s):

["the shade", "a different place"]

Please return a list all of the bridging anaphors in the following text in the order in which they appear. Output the anaphor mention phrase exactly as it appears in the text. If there are no bridging anaphors, return an empty list.

Text:

... making this technique the basis of training for all types of dance . While dancing ballet takes dedication and requires serious training , you can learn the basics to prepare yourself for further study . Learn to get ready for practicing...

Answer(s):

["the basics", "further study"]

Please return a list all of the bridging anaphors in the following text in the order in which they appear. Output the anaphor mention phrase exactly as it appears in the text. If there are no bridging anaphors, return an empty list.

Text:

{text}

Answer(s):

## C.2 Antecedent Selection

You are a linguistic analyst whose job is to select the associative antecedent for of a bridging anaphor: mentions of newly introduced entities (noun phrases) in a text, for which a reader would need to refer back to a previously mentioned, non-identical entity (the antecedent) to resolve their meaning. There are several classes of bridging instances, defined by the associative relationship between the bridging anaphor and its associative antecedent. In the following examples, the bridging antecedent is surrounded by \*asterisks\* and the bridging anaphor is surrounded by {{double curly brackets}}.

comparison-relative: The anaphor is preceded by a comparative marker (other, another, same, more, ordinal modifiers, comparative adjectives, superlatives, etc.) which implies a comparison to the antecedent. For example: "\*The children\* ... {{another child}}" (=another with comparison to the aforementioned children); similar cases may be {{similar children}}, {{older children}}(compared to the aforementioned children), etc.

comparison-sense: the semantic type of a phrase requires a previous mention to identify it, for example "\*the Italian restaurant\* ... {{a Chinese one}}" (we can't know "a Chinese one" is a restaurant without referring back to the Italian restaurant), or "{{another one}}", "{{the others}}" etc.

comparison-time: the anaphor refers to a specific time/timeframe which is understandable with reference to the antecedent, for example: "\*Tuesday, February 2nd\* ... {{the following week}}"

entity-meronymy: the anaphor is a subunit of the antecedent (part-whole), including physical subunits, portion-substance relations, and regions/subsections. For example: "\*the house\* ... {{the door}}" (=of the house).

entity-associative: the anaphor is an attribute or closely associated entity of the antecedent, including both prototypical and inducible associations: "\*a wedding\* ... {{the bride}}" (=the bride at that wedding), implicit arguments of a predicate or a verbal nominalization: "\*a play\*... {{the performance}}" (=of the play), relational nouns: "\*a murder\* ... {{the victim}}"

entity-property: the anaphor is a physical or intangible property of the antecedent (e.g., smell, length, size, style, etc.): "\*the tea\* ... {{the sweet aroma}}"

entity-resultative: the anaphor is logically inferable from the antecedent (e.g., result, transformation/transmutation, cause): "\*the dough\* ... {{the bread}}" (=the dough becomes bread after baking)

set-member: the anaphor is an element of the antecedent set, including groups-member relations and classes-instances: "\*the cars\* ... {{the Mazda}}", additionally indefinite members to definite sets: "\*a candle\*"

on each cupcake ... {{the candles}}"

set-subset: the anaphor is a subset of the antecedent set: "\*the cars\* ... {{the Mazdas}}" (not all Mazdas, just the subset among the aforementioned cars)

set-span-interval: the anaphor is a sub-span of a spatial or temporal interval defined by the antecedent: "\*last week\* ... {{Wednesday}}" (=Wednesday of last week), "\*Sunday\* ... {{the morning}}" (=the morning portion of that Sunday)

other: the anaphor requires a previous entity for interpretation, but it doesn't fit into any of the above categories. This is a rare class.

Here are 2 examples of the task:

Please return a single string for associative antecedent of the bridging anaphor surrounded by {{double curly brackets}}. Output the antecedent mention phrase exactly as it appears in the text. If there is no associative antecedent, return "no antecedent". The antecedent you are returning CANNOT be the same as the bracketed anaphor.

Text:

... with their friends to a picnic. The picnic was supposed to take place in a grove, but {{the shade}} wasn't enough, so they had to find a different place. Conny started to say ...

Answer:

a grove

Please return a single string for associative antecedent of the bridging anaphor surrounded by {{double curly brackets}}. Output the antecedent mention phrase exactly as it appears in the text. If there is no associative antecedent, return "no antecedent".

Text:

... making this technique the basis of training for all types of dance . While dancing ballet takes dedication and requires serious training , you can learn the basics to prepare yourself for {{further study}} . Learn to get ready for practicing ...

Answer:

ballet

Please return a single string for associative antecedent of the bridging anaphor surrounded by {{double curly brackets}}. Output the antecedent mention phrase exactly as it appears in the text. If there is no associative antecedent, return "no antecedent".

Text:

{text}

Answer:

### C.3 Subtype Categorization

You are a linguistic analyst whose job is to select the subtype classification for a bridging anaphor - antecedent pair: mentions of newly introduced entities (the anaphor) in a text, for which a reader would need to refer back to a previously mentioned, non-identical entity (the antecedent) to resolve their meaning. There are several classes of bridging instances, defined by the associative relationship between the bridging anaphor and its associative antecedent. In the following subtype examples, the bridging antecedent is surrounded by \*asterisks\* and the bridging anaphor is surrounded by {{double curly brackets}}.

comparison-relative: The anaphor is preceded by a comparative marker (other, another, same, more, ordinal modifiers, comparative adjectives, superlatives, etc.) which implies a comparison to the antecedent. For example: "\*The children\* ... {{another child}}" (=another with comparison to the aforementioned children); similar cases may be {{similar children}}, {{older children}} (compared to the aforementioned children), etc.

comparison-sense: the semantic type of a phrase requires a previous mention to identify it, for example "\*the Italian restaurant\* ... {{a Chinese one}}" (we can't know "a Chinese one" is a restaurant without referring back to the Italian restaurant), or "{{another one}}", "{{the others}}". etc.

comparison-time: the anaphor refers to a specific time/timeframe which is understandable with reference to the antecedent, for example: "\*Tuesday, February 2nd\* ... {{the following week}}"

entity-meronymy: the anaphor is a subunit of the antecedent (part-whole), including physical subunits, portion-substance relations, and regions/subsections. For example: "\*the house\* ... {{the door}}" (=of the house).

entity-associative: the anaphor is an attribute or closely associated entity of the antecedent, including both prototypical and inducible associations: "\*a wedding\* ... {{the bride}}" (=the bride at that wedding), implicit arguments of a predicate or a verbal nominalization: "\*a play\*... {{the performance}}" (=of the play), relational nouns: "\*a murder\* ... {{the victim}}"

entity-property: the anaphor is a physical or intangible property of the antecedent (e.g., smell, length, size, style, etc.): "\*the tea\* ... {{the sweet aroma}}"

entity-resultative: the anaphor is logically inferable from the antecedent (e.g., result, transformation/transmutation, cause): "\*the dough\* ... {{the bread}}" (=the dough becomes bread after baking)

set-member: the anaphor is an element of the antecedent set, including groups-member relations and classes-instances: "\*the cars\* ... {{the Mazda}}", additionally indefinite members to definite sets: "\*a candle\* on each cupcake ... {{the candles}}"

set-subset: the anaphor is a subset of the antecedent set: "\*the cars\* ... {{the Mazdas}}" (not all Mazdas, just the subset among the aforementioned cars)

set-span-interval: the anaphor is a sub-span of a spatial or temporal interval defined by the antecedent: "\*last week\* ... {{Wednesday}}" (=Wednesday of last week), "\*Sunday\* ... {{the morning}}" (=the morning portion of that Sunday)

other: the anaphor requires a previous entity for interpretation, but it doesn't fit into any of the above categories. This is a rare class.

Here are 2 examples of the task:

In the following text, a bridging anaphora is marked with {{double curly brackets}} and the corresponding antecedent is surrounded by \*asterisks\*. Read the following text and for the bridging anaphor-antecedent pair, classify the variety of bridging subtype relation (defined above) that holds between the two entities. Multiple subtypes may apply to a single pair. Output a string of all applicable subtypes, connected by semicolons (no spaces).

The possible subtype labels are as follows:

comparison-relative  
comparison-sense  
comparison-time  
entity-associative  
entity-meronymy  
entity-property  
entity-resultative  
set-member  
set-subset  
set-span-interval  
other

Antecedent Text:

... with their friends to a picnic. The picnic was supposed to take place in \*a grove\*, but the shade wasn't enough, so they had to find a different place. Conny started to say ...

Anaphor Text:

... to a picnic. The picnic was supposed to take place in a grove, but {{the shade}} wasn't enough, so they had to find a different place. Conny started to say ...

Answer:

entity-associative

In the following text, a bridging anaphora is marked with {{double curly brackets}} and the corresponding antecedent is surrounded by \*asterisks\*. Read the following text and for the bridging anaphor-antecedent pair, classify the variety of bridging subtype relation (defined above) that holds between the two entities. Multiple subtypes may apply to a single pair. Output a string of all applicable subtypes, connected by semicolons (no

spaces).

The possible subtype labels are as follows:

comparison-relative  
comparison-sense  
comparison-time  
entity-associative  
entity-meronymy  
entity-property  
entity-resultative  
set-member  
set-subset  
set-span-interval  
other

Antecedent Text:

... this technique the basis of training for all types of dance . While dancing \*ballet\* takes dedication and requires serious training , you can learn the basics to prepare yourself for further study . Learn to get ready for practicing ...

Anaphor Text:

... making this technique the basis of training for all types of dance . While dancing ballet takes dedication and requires serious training , you can learn the basics to prepare yourself for {{further study}} . Learn to get ready for ...

Answer:

comparison-relative

In the following text, a bridging anaphora is marked with {{double curly brackets}} and the corresponding antecedent is surrounded by \*asterisks\*. Read the following text and for the bridging anaphor-antecedent pair, classify the variety of bridging subtype relation (defined above) that holds between the two entities. Multiple subtypes may apply to a single pair. Output a string of all applicable subtypes, connected by semicolons (no spaces).

The possible subtype labels are as follows:

comparison-relative  
comparison-sense  
comparison-time  
entity-associative  
entity-meronymy  
entity-property  
entity-resultative  
set-member  
set-subset  
set-span-interval  
other

Antecedent Text:

... {antecedent\_text} ...

Anaphor Text:

... {anaphor\_text} ...

Answer:

# Clustering Analysis for Error Detection in Named Entity Recognition Datasets

Matthew Flynn and Timothy Obiso and Sam Newman\* and Constantine Lignos

Michtom School of Computer Science, Brandeis University

{matthewflynn,timothyobiso,lignos}@brandeis.edu

snewman.aa@gmail.com

## Abstract

This paper introduces a method for the automatic detection of annotation errors and corrections in named entity recognition datasets using a novel two-stage dimension reduction of dense sentence embeddings. We first find the top- $n$  principal components of an embedding and then use UMAP for second-stage, non-linear dimension reduction and clustering using different distance metrics. We analyze these clusters using silhouette scores to flag outlier mentions for correction. Using the corrections in the CoNLL# dataset as a benchmark, all of the top-five outliers needed correction, as did 7 of the top-10. This approach also identified 32 of the top-50 outlier mentions that are corrections. This method offers a relatively low-effort way to leverage text embeddings and dimensionality reduction to identify likely annotation errors. We release related code and data at <https://github.com/bltlab/clustering-for-ner>.

## 1 Introduction

Evaluating the quality of named entity recognition (NER) datasets is a labor-intensive process that requires multiple types of expertise. Adjudicators must check whether annotations are correct, whether an entity type forms a coherent category, and even whether the ontology’s categories form semantically coherent groupings. This requires language fluency and, in many cases, domain expertise. Evaluating quality for a single dataset is feasible for a single language, but no single reviewer can assess and standardize across large, multilingual datasets. Thorough manual review becomes impractical.

We propose using silhouette scores on dimension-reduced sentence embeddings to automate the review and evaluation of NER datasets and their labels during the adjudication process or as post-hoc correction. The core intuition is that

mentions of the same type should cluster closer together in embedding space. In these cases, the silhouette scores will be high; outlier mentions will have low or negative silhouette scores. These scores provide a quantitative signal for reviewers to determine edge cases or annotation errors. Silhouette scores act as an analogue for semantic coherence and can speak to how well the defined labels capture the natural semantic groups present in the data. This process can also highlight difficult cases, as a correct location mention that clusters near a person mention may reveal oddities in the annotation guidelines that may confuse annotators or systems.

Our approach uses principal component analysis (PCA, [Hotelling, 1936](#)) to first identify the top- $n$  principal components of a sentence embedding. We then apply uniform manifold approximation (UMAP, [McInnes et al., 2020](#)) to map the principal components into clusters using a variety of distance metrics, including Euclidean, Chebyshev ([Han et al., 2012](#)), and Canberra ([Lance and Williams, 1966](#)), among others, as well as the Bray-Curtis dissimilarity ([Ricotta and Podani, 2017](#)). We compute the silhouette scores on the clusters formed by this dimension reduction pipeline and evaluate the scores to find and analyze the top- $k$  outlier mentions.

We validate this approach with the CoNLL-03 English test set ([Tjong Kim Sang and De Meulder, 2003](#)) and compare outliers against corrections made in CoNLL# ([Rueda et al., 2024](#)). We show that examining the top-10 outliers in the original test set reliably identifies mentions that were later corrected. We also analyze the clusters formed by the corrected test set and demonstrate that the top-50 outliers can also reliably identify mentions that had been corrected.

Our main contributions are as follows. We propose methods for analyzing clusters and outliers of dimension-reduced sentence embeddings

\*Independent researcher. Work completed while at Brandeis University.

for named entity mentions using silhouette scores. We apply these methods on the CoNLL-03 English benchmark dataset using a variety of embedding models. Finally, we introduce Clusters, a command-line utility to facilitate the application and extension of these methods to further datasets and tasks.

## 2 Related Work

There has long been interest in evaluating the quality of datasets in NLP, as some rate of annotation errors is expected and generally accepted. Various methods and approaches to identifying and correcting such mistakes have included manual, semi-automated, and automated approaches. Dickinson (2015) provides an earlier review of such methods.

With the named entity recognition task, much work has focused on the enduring and popular CoNLL-03 English benchmark dataset to identify errors in annotation. Such work has been applied to the test set specifically (Stanislawek et al., 2019; Wang et al., 2019; Rueda et al., 2024), or the entire dataset (Reiss et al., 2020; Muthuraman et al., 2021). These efforts have produced and released corrected versions of the CoNLL-03 English test set, such as ReCoNLL (Fu et al., 2020) and CoNLL# (Rueda et al., 2024).

As the CoNLL-03 data is from the 1996 Reuters Corpus (Lewis et al., 2004), there are other concerns about how this publicly available newswire data may impact the performance of large, modern models that are trained on similar, and potentially the same, data. To evaluate this, the CoNLL++ dataset (Liu and Ritter, 2023) uses the CoNLL-03 annotation guidelines to create a modern version of the dataset using newswire data from 2020. This enables evaluation of how generations of state-of-the-art CoNLL models perform and generalize to modern data of the same domain and format. These authors also release an updated test set of CoNLL-03 English that removes the tabular, ticker-style data, such as sports scores, from the original.

Another popular dataset that has been reviewed is OntoNotes 5.0 (Weischedel, Ralph et al., 2013), which contains 17 different entity types in comparison to the four of CoNLL-03. Bernier-Colborne and Vajjala (2024) review and correct close to 10% of this dataset and observe that these corrections improve performance of models by an average of 1.23% in overall F1-scores, and they note an even larger improvement of more than 10% for certain

Entity Type	Count
LOC	1633
MISC	754
ORG	1701
PER	1594

Table 1: Counts of mentions for each entity type in the corrected CoNLL# English test set

entity types.

Similar work in identifying annotation errors has also been conducted on non-English datasets, such as Uyghur (Abudukelimu et al., 2018), Japanese (Ichihara et al., 2015), and Hindi (Saha et al., 2009).

## 3 Methodology

We implemented an end-to-end pipeline to generate and evaluate clusters for CoNLL-formatted NER datasets. This includes loading and validating the data, generating embeddings, performing dimension reduction, plotting, and reporting.

### 3.1 Dataset

As prior work shows, the dataset from the CoNLL-03 shared task has proven popular over generations of models as a benchmark, and it has been a focus for analysis of annotation errors. The CoNLL-03 ontology consists of the Person (“PER”), Location (“LOC”), Organization (“ORG”), and Miscellaneous (“MISC”) entity types, uses BIO encoding (Ramshaw and Marcus, 1995), and derives its text from the newswire domain (Lewis et al., 2004). We analyze the outliers identified in both the original English test set and in the corrected CoNLL# test set, which was chosen as it follows the original annotation guidelines when making corrections. Entity type counts for the mentions included in the CoNLL# test set are given in Table 1.

As CoNLL# also corrects sentence boundaries that were incorrectly split in the original English test set, it was necessary to align the original version to the corrected test set to ensure the correct mapping of mentions across datasets for analysis. This includes correcting mentions that were split across the original sentence boundaries.

### 3.2 Loading and Validating Data

To ensure data integrity and consistency, all data was loaded and validated using SeqScore (Palen-Michel et al., 2021; Lignos et al., 2023), an evaluation and validation toolkit for NER. The list of tokens for each mention was joined into a single

string for embedding, and it was mapped to its entity type. For example, a mention sequence with labels [ 'B-PER', 'I-PER', 'I-PER' ] was mapped to PER for analysis and plotting.

### 3.3 Embedding Mentions

We selected a series of freely-available and self-hostable embedding models using the best scores on the Massive Multilingual Text Embedding Benchmark (MMTEB, Enevoldsen et al., 2025).<sup>1</sup> We also selected models based on foundational encoder models, such as SBERT (Reimers and Gurevych, 2019) and XLM-RoBERTa (Conneau et al., 2020). These models vary in architectures, context size, and the dimensions of their output embeddings.

For embedding, we used Qwen’s Qwen3 embedding model family (Zhang et al., 2025),<sup>2</sup> Tencent’s Gemma 3-based KaLM embedding model (Hu et al., 2025; Zhao et al., 2025),<sup>3</sup> intfloat’s XLM-RoBERTa-based Multilingual E5 Instruct (Wang et al., 2024),<sup>4</sup> and the SBERT-based all-MiniLM-L6-v2.<sup>5</sup>

Our implementation supports any embedding model that can be reached at an OpenAI API-compatible embedding endpoint. We downloaded all models from HuggingFace and self-hosted them with vLLM (Kwon et al., 2023) to efficiently generate the embeddings and maintain a consistent embedding interface across models.

### 3.4 Prompts

All of the embedding models, with the exception of all-MiniLM-L6-v2, are instruction-tuned and expect a prompt as part of the embedding input. To standardize this input parameter, we used the same prompt template for all models, including all-MiniLM-L6-v2. This template includes a prompt instruction to inform the model of the purpose of the task, the sentence containing the mention as context, and the mention itself separately. We provide our full prompt template in the Appendix A.1.

### 3.5 Dimension Reduction

We reduced the dimensions of all embeddings before clustering and analysis to avoid the curse of

dimensionality (Peng et al., 2025). We first found the top- $n$  principal components of each embedding with scikit-learn’s (Pedregosa et al., 2011) implementation of PCA. We then used UMAP to generate clusters from the  $n$ -dimension principal components using a variety of distance metrics, including Canberra (Lance and Williams, 1966), Correlation (Székely et al., 2007), Chebyshev, Cosine, Euclidean, Manhattan, and Minkowski (Han et al., 2012), in addition to the Bray-Curtis dissimilarity (Ricotta and Podani, 2017).

### 3.6 Silhouette Scores

With the clusters generated from the dimension-reduced embeddings, the silhouette score for each mention is calculated. This score is a metric for evaluating clusters, where each point’s silhouette score  $s(i)$  is the difference between its average distance from points in the next-nearest cluster  $b(i)$ , and its average distance from points in its own cluster  $a(i)$ . This difference is divided by the max of either  $a(i)$  or  $b(i)$  to obtain the respective point’s silhouette score. This normalizes silhouette scores in the range  $[-1, 1]$ . A higher score implies better clustering (Shahapure and Nicholas, 2020), and, in general, a score  $s(i) > 0.7$  signals strong clustering, and a score  $0.7 > s(i) > 0.5$  is reasonable. A score close to 0.0 represents overlapping clustering.

$$a(i) = \frac{1}{|C(i)| - 1} \sum_{\substack{j \in C(i) \\ j \neq i}} d(i, j), \quad b(i) = \min_{C \neq C(i)} \frac{1}{|C|} \sum_{j \in C} d(i, j),$$

$$s(i) = \frac{b(i) - a(i)}{\max\{a(i), b(i)\}}, \quad S = \frac{1}{N} \sum_{i=1}^N s(i). \quad (1)$$

We used these silhouette scores to identify outlier mentions for analysis.

## 4 Experiments

We completed a comprehensive grid search across all models summarized in Table 2 for both of our experiments. Our first experiment identifies mentions in the original CoNLL-03 English test set that were later corrected in CoNLL#, and the second identifies mentions in CoNLL# that are corrections.

For each experiment, we evaluate the top-performing models and configurations, the relative performance of the three different-sized embedding models in the Qwen3 family, and the relative performance of all models and architectures.

<sup>1</sup>[huggingface.co/spaces/mteb/leaderboard](https://huggingface.co/spaces/mteb/leaderboard)

<sup>2</sup>[huggingface.co/collections/Qwen/qwen3-embedding](https://huggingface.co/collections/Qwen/qwen3-embedding)

<sup>3</sup>[huggingface.co/tencent/KaLM-Embedding-Gemma3-12B-2511](https://huggingface.co/tencent/KaLM-Embedding-Gemma3-12B-2511)

<sup>4</sup>[huggingface.co/intfloat/multilingual-e5-large-instruct](https://huggingface.co/intfloat/multilingual-e5-large-instruct)

<sup>5</sup>[huggingface.co/sentence-transformers/all-MiniLM-L6-v2](https://huggingface.co/sentence-transformers/all-MiniLM-L6-v2)

Model	Parameters	Dimensions
Qwen/Qwen3-Embedding-0.6B	0.6B	1024
Qwen/Qwen3-Embedding-4B	4B	2048
Qwen/Qwen3-Embedding-8B	8B	4096
tencent/KaLM-Embedding-Gemma3-12B-2511	11.76B	3840
intfloat/multilingual-e5-large-instruct	0.6B	512
sentence-transformers/all-MiniLM-L6-v2	22.7M	384

Table 2: Comparison of embedding model sizes (parameter count) and embedding dimensions

Hyperparameter	Values
PCA Components	50, 75, 100
Clustering Components	16, 32, 64
Distance Metric	Bray-Curtis, Canberra, Chebyshev, Correlation, Cosine, Euclidean, Manhattan, Minkowski

Table 3: Hyperparameters for the grid search to find the optimal configuration for each model to identify mentions that were corrected

#### 4.1 Identifying Mentions that were Corrected

We evaluate the performance of models at identifying how many of the top-five and top-10 of their outlier mentions were later corrected by CoNLL#. For this experiment, we held the number of UMAP cluster neighbors constant at 100, and the UMAP clustering used labeled data when learning the clusters. The hyperparameters for this experiment are summarized in Table 3.

#### 4.2 Identifying Mentions that are Corrections

We evaluate performance by comparing how many of the top-50 outlier mentions identified by the model are corrections. As before, we held the number of UMAP cluster neighbors constant at 100; however, we varied the PCA implementation across the original, truncated, and kernel versions of PCA. We also varied the cluster learning with and without labels. The hyperparameters for this experiment are summarized in Table 4.

### 5 Results

Our experiments provide insight into optimal dimension reduction ratios, model architectures, embedding sizes, hyperparameter configurations, and distance metrics.

#### 5.1 Identifying Mentions that were Corrected

We now report model and configuration performance on identifying mentions that were corrected

Hyperparameter	Values
PCA Implementation	Original, Kernel, Truncated
PCA Components	50, 75, 100
Clustering Components	16, 32, 64
Cluster Labels	True, False
Distance Metric	Bray-Curtis, Canberra, Chebyshev, Correlation, Cosine, Euclidean, Manhattan, Minkowski

Table 4: Hyperparameters for the grid search to find the optimal configuration for each model to identify mentions that are corrections

in CoNLL#.

#### 5.1.1 Top-Performing Models

Among its top-five and top-10 outliers, our best-performing model correctly identified five and seven mentions, respectively, that were later corrected. The three mentions that were not correctly identified as later corrected are all MISC mentions. Three different models reported among the top-five scores, with four of the top-five scores using Manhattan distance when learning clusters with UMAP. For the fourth and fifth best runs, using the KaLM-Embedding-Gemma3-12B-2511 model, the hyperparameters only differed in their distance metric. Interestingly, Canberra can be interpreted as a weighted Manhattan distance, and for this model, it performed slightly worse than its unweighted Manhattan counterpart. These results are summarized in Table 5, and 2d and 3d t-SNE (Hinton and Roweis, 2002) and UMAP projections are available in Appendix A.2.<sup>6</sup>

<sup>6</sup>While the authors were not aware of this at the time of submission, Peter Mayhew blogged about using 2d t-SNE projections of mention embeddings to help interpret tags and mentions from the CoNLL-03 English dataset. We refer the reader to his blog for more discussion: [mayhewsw.github.io/2022/01/30/conll-span-embeddings/](https://mayhewsw.github.io/2022/01/30/conll-span-embeddings/)

Model	Metric	PCA Comp	Cluster Comp	Top-5	Top-10
Qwen3-Embedding-4B	Manhattan	50	32	5	7
multilingual-e5-large-instruct	Manhattan	50	32	5	5
multilingual-e5-large-instruct	Manhattan	50	16	5	5
KaLM-Embedding-Gemma3-12B-2511	Manhattan	50	16	4	6
KaLM-Embedding-Gemma3-12B-2511	Canberra	50	16	4	4

Table 5: Best-performing model configurations at identifying mentions that were later corrected

Entity Type	Mention	Score	CoNLL#	Corrected
ORG	Portsmouth	-0.9270	LOC	✓
ORG	Oxford	-0.9269	LOC	✓
LOC	DURBAN	-0.9235	PER	✓
LOC	SANTIAGO	-0.9218	PER	✓
ORG	DENVER	-0.9139	MISC	✓
ORG	GREEN BAY	-0.9041	MISC	✓
MISC	Lombardi Award	-0.9022	MISC	
ORG	OHIO STATE	-0.9021	MISC	✓
MISC	LOMBARDI AWARD	-0.9018	MISC	
MISC	AMERICAN	-0.9015	MISC	

Table 6: Top-10 outliers as identified by the top-performing Qwen3-Embedding-4B model

### 5.1.2 Qwen3 Embedding Model Family

The three different Qwen3 embedding models provide an interesting opportunity to compare the relative performance of the different sizes of this model family on the same task and using the same grid search to identify the optimal model configurations. Table 7 summarizes the Qwen3 scores for the first experiment.

Interestingly, the 4B model with 2048 embedding dimensions performed best among the Qwen3 embedding family, and the best overall for this experiment.

This finding runs counter to the intuition that the largest model with the most embedding dimensions should be able to perform better than a smaller counterpart, as is often the case in MMTEB performance. However, with the dimension-reduced embeddings, the 4B model performed best. Additionally, the optimal configuration for each model used a different distance metric. These metrics are related, however, as the Manhattan distance used by the 4B model is a Minkowski distance where  $p = 1$ , with the 8B model performing best with the default UMAP Minkowski configuration where  $p = 2$ . As noted above, the Canberra distance used by the 0.6B model can be interpreted as a weighted Manhattan distance. These results are summarized in Table 7.

### 5.1.3 All Architectures

We also evaluated the relative performance of each model family and architecture on this task, and the

top score for each model is summarized in Table 8.

Consistent with other models, the best configuration for all-MiniLM-L6-v2, the only remaining unreported model for this first experiment, also reported its best run using Manhattan distance. This result also follows the intuition that the smallest, non-instruction-tuned embedding model would perform worse relative to the other larger, instruction-tuned embedding models.

## 5.2 Identifying Mentions that are Corrections

We now report model and configuration performance on identifying mentions that are corrections in CoNLL#, which includes some corrections made in prior attempts to correct the CoNLL-03 English data.

### 5.2.1 Top-Performing Models

Unlike the results in the first experiment, the best-performing model for identifying mentions in CoNLL# that are corrections is the small, non-instruction-tuned all-MiniLM-L6-v2. It reports all five of the top-five runs, and its top-scoring run correctly identifies 32 corrected mentions among its top-50 outliers. Similar to experiment one, Canberra distance also performs well at this task, as well as Chebyshev. These results are summarized in Table 9, and 2d and 3d t-SNE and UMAP projections are available in Appendix A.3.

Table 10 summarizes the top-10 outliers identified by all-MiniLM-L6-v2. All are MISC mentions that were corrected from either LOC or O mentions

Model	Metric	PCA Comp	Cluster Comp	Top-5	Top-10
Qwen3-Embedding-0.6B	Canberra	50	32	4	4
Qwen3-Embedding-4B	Manhattan	50	32	5	7
Qwen3-Embedding-8B	Minkowski	50	16	4	6

Table 7: Performance of Qwen3 embedding family at identifying mentions that were later corrected

Model	Metric	PCA Comp	Cluster Comp	Top-5	Top-10
Qwen3-Embedding-4B	Manhattan	50	32	5	7
multilingual-e5-large-instruct	Manhattan	50	32	5	5
KaLM-Embedding-Gemma3-12B-2511	Manhattan	50	16	4	6
Qwen3-Embedding-8B	Minkowski	50	16	4	6
Qwen3-Embedding-0.6B	Canberra	50	32	4	4
all-MiniLM-L6-v2	Manhattan	50	16	3	3

Table 8: Performance of each embedding model family and architecture at identifying mentions that were later corrected

in CoNLL#. This pattern suggests that MISC is the noisiest of the labels in the CoNLL-03 ontology and that, while all of these mentions form a cluster given their similar or identical textual content, they are distant from the core MISC cluster. Notably, these MISC mentions have no context, as they are entire sentences in the CoNLL dataset and exist among the ticker-style sports scores present in the test set.

### 5.2.2 Qwen3 Embedding Model Family

As before, we compared the relative performance of the three sizes of Qwen3 embedding. Unlike the first experiment, the largest model with 8B parameters significantly outperformed the smaller models. It neared the performance of all-MiniLM-L6-v2, identifying 30 corrected mentions among its top-50 outliers. These results are summarized in Table 11.

All three Qwen models performed best with the Bray-Curtis dissimilarity metric, which differs from the classic or weighted Minkowski distance metrics that score best for other models. Bray-Curtis is a popular statistical metric in ecology and biology, and it quantifies the difference between two different samples.

### 5.2.3 All Architectures

These results show that the best-performing model overall is all-MiniLM-L6-v2, which is also the smallest of the models we evaluated. Conversely, KaLM-Embedding-Gemma3-12B-2511 is the largest of the models we evaluated and while it boasts the strongest performance on MMTEB, it performed significantly worse than all other models across any configuration for this second experiment. It identified just five corrected mentions among its

top-50 outliers. These results are summarized in Table 12.

## 6 Discussion

This method provides insight into automated detection of annotation errors in NER, as well as the relative performance of model sizes and architectures in the clustering of mentions. It can also be useful in designing ontologies and creating datasets.

### 6.1 Identifying Annotation Errors

Our implementation lists the top- $k$  outlier mentions for each entity type in a dataset and overall. These outliers are useful for surfacing potential annotation errors for closer review without needing to review each mention in a dataset individually, and for evaluating a corrected dataset to ensure corrected mentions are consistent with existing annotations.

Treating CoNLL# as the gold standard, our experiments show that our method can identify among its top-five and top-10 outliers, five and seven labels that were corrected. Likewise, the same pipeline when applied to the gold CoNLL# data can identify 32 mentions that are corrections among its top-50 outliers.

Our experiments also reveal that different model sizes and architectures perform differently on this task. In identifying mentions that were corrected, larger models showed superior performance using weighted and unweighted Minkowski distance metrics where  $p = 1$  or  $p = 2$ . However, at identifying corrected mentions, the smallest model, all-MiniLM-L6-v2 with only 22.7M parameters and embeddings with 384 dimensions, performed best. This finding makes it possible to use this method

Model	Metric	Labels	PCA Implementation	PCA Comp	Cluster Comp	Corrected
all-MiniLM-L6-v2	Canberra	False	PCA	50	32	32
all-MiniLM-L6-v2	Chebyshev	False	Truncated	50	16	31
all-MiniLM-L6-v2	Canberra	False	Kernel	50	32	31
all-MiniLM-L6-v2	Chebyshev	False	PCA	50	32	31
all-MiniLM-L6-v2	Chebyshev	False	PCA	50	64	31

Table 9: Best-performing model configurations at identifying mentions that are corrections

Entity Type	Mention	Score	CoNLL#	Corrected
LOC	ATLANTIC DIVISION	-0.2328	MISC	✓
LOC	ATLANTIC DIVISION	-0.2317	MISC	✓
LOC	PACIFIC DIVISION	-0.2311	MISC	✓
LOC	PACIFIC DIVISION	-0.2310	MISC	✓
O	CENTRAL DIVISION	-0.2299	MISC	✓
O	CENTRAL DIVISION	-0.2288	MISC	✓
O	CENTRAL DIVISION	-0.2274	MISC	✓
O	EASTERN DIVISION	-0.2239	MISC	✓
O	EASTERN DIVISION	-0.2238	MISC	✓
O	WESTERN DIVISION	-0.2238	MISC	✓

Table 10: Top-10 outliers as identified by all-MiniLM-L6-v2

with limited computing resources, and even on just a CPU. Its top scores were also with Canberra and Chebyshev distance, while the large, instruction-tuned pooling models performed best with the Bray-Curtis dissimilarity metric at this task.

## 6.2 Dataset Creation

Dataset creation can be a time-consuming and burdensome process, especially in domain-specific, multilingual, or low-resourced settings, where finding domain experts or native speakers available for quality data annotation work can be difficult, and potentially expensive. Using the top- $k$  outlier mentions for each entity type during the annotation workflow can reveal errors and tough mentions. These tough mentions can also provide insight and data points for any labels that may be unclear or borderline, which can help improve and clarify annotation guidelines and the ontology.

Strategies, such as the MATTER lifecycle (Pustejovsky and Stubbs, 2012) or some other workflow, exist for this process. These strategies include ontology creation and review, in addition to adjudication and annotation, as part of the overall dataset creation process. Data annotation tools, such as brat (Stenetorp et al., 2012) and Label Studio (Tkachenko et al., 2020-2025), among others, also exist to facilitate and streamline this process. Integrating our method as part of the workflow can reduce burden as it quickly surfaces the top- $k$  outliers for each label at any step. With this quantitative information and supporting visual projections,

dataset authors can improve the quality of their annotated datasets and provide consistent feedback to annotators during the annotation process as part of a continuous integration workflow and without creating additional burdens for reviewers.

## 6.3 MISC Mentions

Being a catchall entity type, it follows that MISC mentions are the noisiest and least semantically coherent. All of the missed mentions in our first experiment were MISC, and our second experiment shows that among the top-10 outlier mentions identified by the top-performing all-MiniLM-L6-v2 model, all were very similar, and sometimes identical, MISC mentions. All were corrected labels as well. This suggests that MISC mentions were indeed the noisiest and least coherent.

Our method can quickly surface all outlier MISC mentions, report silhouette scores for all their mentions for comparison, reveal their relative coherence given the MISC type’s catchall definition, and quantitatively show the relative value of MISC in an ontology. Our method can help ontology designers decide if it is better to split MISC into other entity types, or to drop it altogether.

## 7 Future Work and Conclusion

We show that using silhouette scores and dimension-reduced embeddings to evaluate ontologies and datasets can automatically identify annotation errors in original datasets and corrected mentions in updated datasets. To facilitate such

Model	Metric	Labels	PCAI	PCAC	CC	Corrected
Qwen3-Embedding-0.6B	Bray-Curtis	False	PCA	75	16	12
Qwen3-Embedding-4B	Bray-Curtis	False	Kernel	75	16	24
Qwen3-Embedding-8B	Bray-Curtis	False	PCA	50	16	<b>30</b>

Table 11: Performance of Qwen3 embedding family at identifying mentions that are corrections

Model	Metric	Labels	PCA Impl	PCA Comp	Cluster Comp	Corr
all-MiniLM-L6-v2	Canberra	False	PCA	50	32	32
Qwen3-Embedding-8B	Bray-Curtis	False	Kernel	100	16	30
multilingual-e5-large-instruct	Chebyshev	False	Truncated	50	32	28
Qwen3-Embedding-4B	Bray-Curtis	False	Kernel	75	16	24
Qwen3-Embedding-0.6B	Bray-Curtis	False	PCA	75	16	12
KaLM-Embedding-Gemma3-12B-2511	Bray-Curtis	False	PCA	75	32	5

Table 12: Performance of all architectures at identifying mentions that are corrections

future work, we release the Clusters utility under a permissive open source license.

Future work can extend this approach and the Clusters utility to other NLP domains that use labeled data. It is also possible to extend the two-stage dimension reduction to unlabeled data, such as that used to train LLMs, in order to analyze the semantic distribution of the training data, and to help ensure it is balanced and not biased to any domain or topic. The approach can also be applied for analysis in any of the numerous domains that use embeddings, such as dense information retrieval.

## Limitations

With dimension reduction using PCA and UMAP, there is inherently some amount of information loss. Depending on the task or embedding size, that may be undesirable and would limit the applicability of our method.

Due to the limited public availability of previously corrected NER datasets, we were forced to confine our experiments to English data.

Modern language models are often trained on large and diverse corpora, with popular encoding models trained on 100 languages (Conneau et al., 2020) and some modern LLMs trained on 140+ languages (Team et al., 2025). The issue of low-resourced languages, however, remains. This clustering analysis is ultimately reliant on the performance of embedding models to generate the embeddings that are used for the silhouette score analysis. If the encoding model does not perform well or understand a low-resourced language corpus, the performance of this approach may be less reliable.

Likewise, embedding models and LLMs alike inherit the biases of their training data, and that

can also impact the usefulness of their generated embeddings at this task.

## Acknowledgments

Constantine Lignos was partially supported by the grant *Improving Relevance and Recovery by Extracting Latent Query Structure* from eBay to Brandeis University.

## References

- Halidanmu Abudukelimu, Abudoukelimu Abulizi, Boliang Zhang, Xiaoman Pan, Di Lu, Heng Ji, and Yang Liu. 2018. [Error analysis of Uyghur name tagging: Language-specific techniques and remaining challenges](#). In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan. European Language Resources Association (ELRA).
- Gabriel Bernier-Colborne and Sowmya Vajjala. 2024. [Annotation Errors and NER: A Study with OntoNotes 5.0](#). *Preprint*, arXiv:2406.19172.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. [Unsupervised Cross-lingual Representation Learning at Scale](#). *Preprint*, arXiv:1911.02116.
- Markus Dickinson. 2015. [Detection of Annotation Errors in Corpora](#). *Language and Linguistics Compass*, 9(3):119–138.
- Kenneth Enevoldsen, Isaac Chung, Imene Kerboua, Márton Kardos, Ashwin Mathur, David Stap, Jay Gala, Wissam Siblini, Dominik Krzemiński, Genta Indra Winata, Saba Sturua, Saiteja Utpala, Mathieu Ciancone, Marion Schaeffer, Gabriel Sequeira, Diganta Misra, Shreeya Dhakal, Jonathan Ryrstrøm, Roman Solomatin, and 67 others. 2025.

- MMTEB: Massive Multilingual Text Embedding Benchmark. *Preprint*, arXiv:2502.13595.
- Jinlan Fu, Pengfei Liu, and Qi Zhang. 2020. **Rethinking Generalization of Neural Models: A Named Entity Recognition Case Study**. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(05):7732–7739.
- Jiawei Han, Micheline Kamber, and Jian Pei. 2012. **2 - getting to know your data**. In *Data Mining: Concepts and Techniques (Third Edition)*, third edition edition, The Morgan Kaufmann Series in Data Management Systems, pages 39–82. Morgan Kaufmann, Boston.
- Geoffrey E Hinton and Sam Roweis. 2002. **Stochastic Neighbor Embedding**. In *Advances in Neural Information Processing Systems*, volume 15. MIT Press.
- Harold Hotelling. 1936. Relations Between Two Sets of Variates. *Biometrika*, 28(3/4):321–377.
- Xinshuo Hu, Zifei Shan, Xinpeng Zhao, Zetian Sun, Zhenyu Liu, Dongfang Li, Shaolin Ye, Xinyuan Wei, Qian Chen, Baotian Hu, Haofen Wang, Jun Yu, and Min Zhang. 2025. **KaLM-Embedding: Superior Training Data Brings A Stronger Embedding Model**. *Preprint*, arXiv:2501.01028.
- Masaaki Ichihara, Kanako Komiya, Tomoya Iwakura, and Maiko Yamazaki. 2015. Error analysis of named entity recognition in BCCWJ. *Recall*, 61:2641.
- Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph E. Gonzalez, Hao Zhang, and Ion Stoica. 2023. Efficient Memory Management for Large Language Model Serving with PagedAttention. In *Proceedings of the ACM SIGOPS 29th Symposium on Operating Systems Principles*.
- Godfrey N Lance and William T Williams. 1966. Computer programs for hierarchical polythetic classification (“similarity analyses”). *The Computer Journal*, 9(1):60–64.
- David D. Lewis, Yiming Yang, Tony G. Rose, and Fan Li. 2004. RCV1: A New Benchmark Collection for Text Categorization Research. *J. Mach. Learn. Res.*, 5:361–397.
- Constantine Lignos, Maya Kruse, and Andrew Rueda. 2023. **Improving NER research workflows with SeqScore**. In *Proceedings of the 3rd Workshop for Natural Language Processing Open Source Software (NLP-OSS 2023)*, pages 147–152, Singapore. Association for Computational Linguistics.
- Shuheng Liu and Alan Ritter. 2023. **Do CoNLL-2003 named entity taggers still work well in 2023?** In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8254–8271, Toronto, Canada. Association for Computational Linguistics.
- Leland McInnes, John Healy, and James Melville. 2020. **UMAP: Uniform Manifold Approximation and Projection for Dimension Reduction**. *Preprint*, arXiv:1802.03426.
- Karthik Muthuraman, Frederick Reiss, Hong Xu, Bryan Cutler, and Zachary Eichenberger. 2021. **Data cleaning tools for token classification tasks**. In *Proceedings of the Second Workshop on Data Science with Human in the Loop: Language Advances*, pages 59–61, Online. Association for Computational Linguistics.
- Chester Palen-Michel, Nolan Holley, and Constantine Lignos. 2021. **SeqScore: Addressing barriers to reproducible named entity recognition evaluation**. In *Proceedings of the 2nd Workshop on Evaluation and Comparison of NLP Systems*, pages 40–50, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. 2011. Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research*, 12:2825–2830.
- Dehua Peng, Zhipeng Gui, and Huayi Wu. 2025. **Interpreting the curse of dimensionality from distance concentration and manifold effect**. *Preprint*, arXiv:2401.00422.
- James Pustejovsky and Amber Stubbs. 2012. *Natural Language Annotation for Machine Learning: A guide to corpus-building for applications*. " O’Reilly Media, Inc." .
- Lance Ramshaw and Mitch Marcus. 1995. **Text chunking using transformation-based learning**. In *Third Workshop on Very Large Corpora*.
- Nils Reimers and Iryna Gurevych. 2019. **Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks**. *Preprint*, arXiv:1908.10084.
- Frederick Reiss, Hong Xu, Bryan Cutler, Karthik Muthuraman, and Zachary Eichenberger. 2020. **Identifying incorrect labels in the CoNLL-2003 corpus**. In *Proceedings of the 24th Conference on Computational Natural Language Learning*, pages 215–226, Online. Association for Computational Linguistics.
- C. Ricotta and J. Podani. 2017. **On some properties of the Bray-Curtis dissimilarity and their ecological meaning**. *Ecological Complexity*, 31:201–205.
- Andrew Rueda, Elena Alvarez-Mellado, and Constantine Lignos. 2024. **CoNLL#: Fine-grained error analysis and a corrected test set for CoNLL-03 English**. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 3718–3728, Torino, Italia. ELRA and ICCL.

- Sujan Kumar Saha, Sudeshna Sarkar, and Pabitra Mitra. 2009. [Hindi named entity annotation error detection and correction](#). *Language Forum*, 35(2):73–93. Publisher: Bahri Publications.
- Ketan Rajshekhar Shahapure and Charles Nicholas. 2020. [Cluster Quality Analysis Using Silhouette Score](#). In *2020 IEEE 7th International Conference on Data Science and Advanced Analytics (DSAA)*, pages 747–748.
- Tomasz Stanislawek, Anna Wróblewska, Alicja Wójcicka, Daniel Ziembicki, and Przemyslaw Biecek. 2019. [Named entity recognition - is there a glass ceiling?](#) In *Proceedings of the 23rd Conference on Computational Natural Language Learning (CoNLL)*, pages 624–633, Hong Kong, China. Association for Computational Linguistics.
- Pontus Stenetorp, Sampo Pyysalo, Goran Topić, Tomoko Ohta, Sophia Ananiadou, and Jun’ichi Tsujii. 2012. brat: a Web-based Tool for NLP-Assisted Text Annotation. In *Proceedings of the Demonstrations Session at EACL 2012*, Avignon, France. Association for Computational Linguistics.
- Gábor J. Székely, Maria L. Rizzo, and Nail K. Bakirov. 2007. [Measuring and testing dependence by correlation of distances](#). *The Annals of Statistics*, 35(6).
- Gemma Team, Aishwarya Kamath, Johan Ferret, Shreya Pathak, Nino Vieillard, Ramona Merhej, Sarah Perrin, Tatiana Matejovicova, Alexandre Ramé, Morgane Rivière, Louis Rouillard, Thomas Mesnard, Geoffrey Cideron, Jean-Bastien Grill, Sabela Ramos, Edouard Yvinec, Michelle Casbon, Etienne Pot, Ivo Penchev, and 197 others. 2025. [Gemma 3 Technical Report](#). Preprint, arXiv:2503.19786.
- Erik F. Tjong Kim Sang and Fien De Meulder. 2003. [Introduction to the CoNLL-2003 shared task: Language-independent named entity recognition](#). In *Proceedings of the Seventh Conference on Natural Language Learning at HLT-NAACL 2003*, pages 142–147.
- Maxim Tkachenko, Mikhail Malyuk, Andrey Holmanyuk, and Nikolai Liubimov. 2020–2025. [Label Studio: Data labeling software](#). Open source software available from <https://github.com/HumanSignal/label-studio>.
- Liang Wang, Nan Yang, Xiaolong Huang, Linjun Yang, Rangan Majumder, and Furu Wei. 2024. Multilingual E5 Text Embeddings: A Technical Report. *arXiv preprint arXiv:2402.05672*.
- Zihan Wang, Jingbo Shang, Liyuan Liu, Lihao Lu, Jiacheng Liu, and Jiawei Han. 2019. [CrossWeigh: Training named entity tagger from imperfect annotations](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 5154–5163, Hong Kong, China. Association for Computational Linguistics.
- Weischedel, Ralph, Palmer, Martha, Marcus, Mitchell, Hovy, Eduard, Pradhan, Sameer, Ramshaw, Lance, Xue, Nianwen, Taylor, Ann, Kaufman, Jeff, Franchini, Michelle, El-Bachouti, Mohammed, Belvin, Robert, and Houston, Ann. 2013. [OntoNotes Release 5.0](#).
- Yanzhao Zhang, Mingxin Li, Dingkun Long, Xin Zhang, Huan Lin, Baosong Yang, Pengjun Xie, An Yang, Dayiheng Liu, Junyang Lin, Fei Huang, and Jingren Zhou. 2025. Qwen3 Embedding: Advancing Text Embedding and Reranking Through Foundation Models. *arXiv preprint arXiv:2506.05176*.
- Xinping Zhao, Xinshuo Hu, Zifei Shan, Shouzheng Huang, Yao Zhou, Xin Zhang, Zetian Sun, Zhenyu Liu, Dongfang Li, Xinyuan Wei, Youcheng Pan, Yang Xiang, Meishan Zhang, Haofen Wang, Jun Yu, Baotian Hu, and Min Zhang. 2025. [KaLM-Embedding-v2: Superior Training Techniques and Data Inspire A Versatile Embedding Model](#). Preprint, arXiv:2506.20923.

## A Appendix

### A.1 Prompts

PROMPT: Embed the following for use in clustering analysis with dimension reduction

```
### Context
{context}
```

### Mention

```
{mention}
```

In this template, "context" represents the entire sentence that contains the mention, and "mention" represents the mention itself.

## A.2 Projections of Mentions that were Corrected

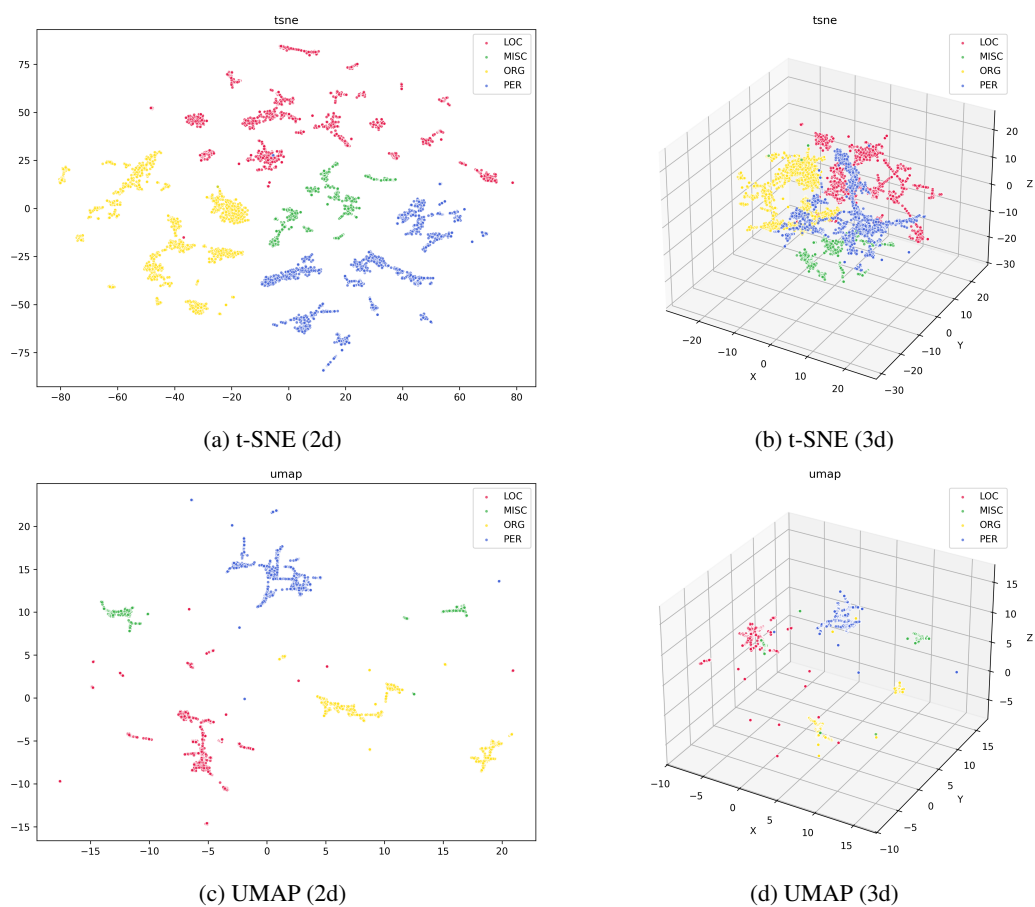


Figure 1: t-SNE and UMAP projections in 2d and 3d for the best-performing run with Qwen3-Embedding-4B at identifying mentions that were corrected

### A.3 Projections of Mentions that are Corrections

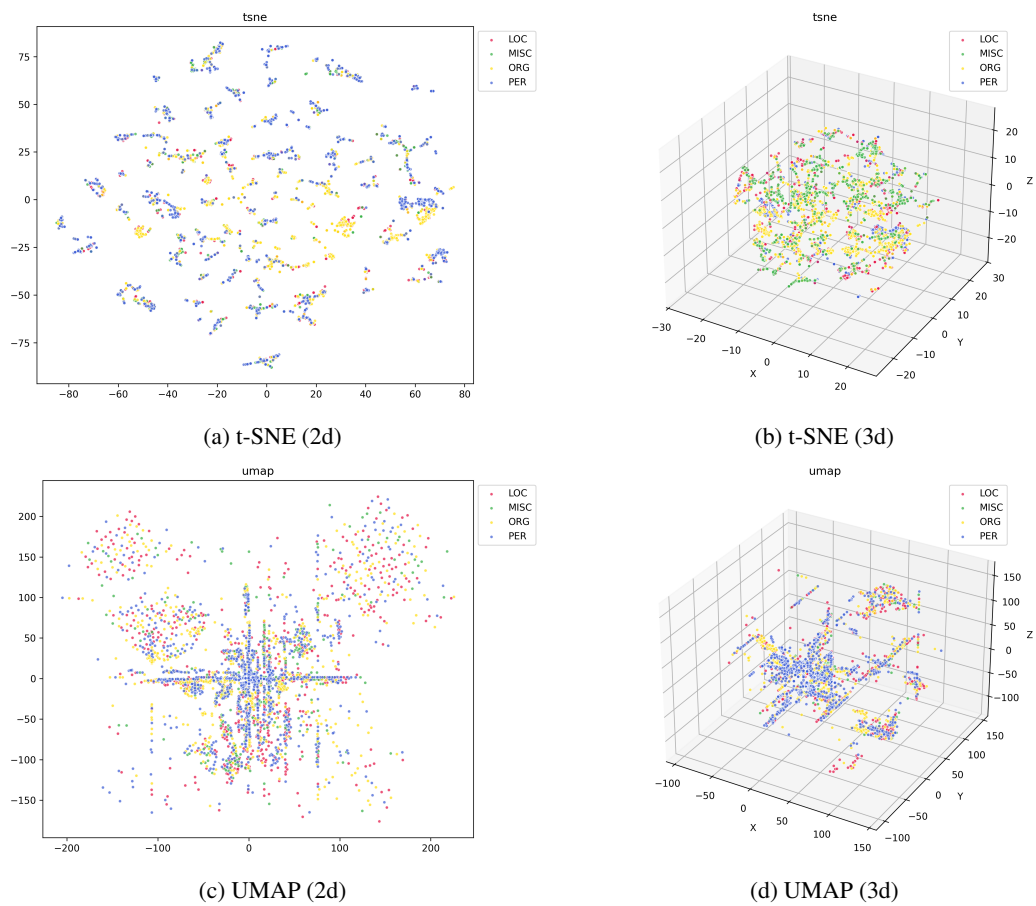


Figure 2: t-SNE and UMAP projections in 2d and 3d for the best-performing run with all-MiniLM-L6-v2 at identifying mentions that are corrections

# When Ground Truth Disagrees: A Human-in-the-Loop Audit of Annotation Errors in High-Stakes Crash Narratives

Md Sajjad Hossain<sup>1</sup>, Lin Li<sup>1</sup>, Judy A. Perkins<sup>2</sup>, John Clary<sup>3</sup>, Joel Meyer<sup>3</sup>

<sup>1</sup>Department of Computer Science    <sup>2</sup>Department of Civil & Environmental Engineering  
Prairie View A&M University, Prairie View, Texas

<sup>3</sup>Austin Transportation & Public Works, Austin, Texas

mhossain2@pvamu.edu, lilin@pvamu.edu, juperkins@pvamu.edu  
{John.Clary, joel.meyer}@austintexas.gov

## Abstract

Linguistic annotation of high-stakes narrative data is often constrained by data confidentiality, domain expertise, and the lack of large-scale multi-annotator pipelines. We present a human-in-the-loop framework for auditing annotation discrepancies in crash narratives, combining structured labels, narrative-based annotation, and expert adjudication. Using 9,387 crash reports, we conduct a multi-layer analysis of disagreement across annotation sources. Nearly half of the records (49.4%) exhibit discrepancies between structured and narrative labels, driven mainly by unsupported structured assignments. In contrast, narrative-based annotation achieves near-perfect agreement with adjudication ( $\kappa = 0.990$ ), indicating strong consistency when grounded in textual evidence. We introduce a taxonomy of discrepancies, showing refinement opportunities and missing details are the most common, while linguistic factors such as hedging and under-specification contribute to ambiguity. We further show that annotator-reported uncertainty strongly predicts annotation difficulty, with uncertain records nearly nine times more likely to disagree with structured labels. These findings highlight limitations of administrative coding and support a scalable, uncertainty-guided annotation paradigm for restricted-access domains.

## 1 Introduction

Linguistic annotation of natural language corpora is a foundational component of modern NLP systems, enabling both supervised learning and evaluation of language understanding models. In high-stakes domains such as transportation safety, annotation quality is particularly critical (Klie et al., 2024; Kumar and Sangwan, 2025; Khasanah et al., 2025), as labeled data directly informs policy decisions, risk analysis, and system design. However, annotation in such settings is often constrained by

data confidentiality, limited access to domain expertise, and reliance on administrative coding systems, which may not fully capture the underlying textual evidence (Di Martino et al., 2020; Salami, 2023; Campbell and Giadresco, 2020).

Crash reporting provides a representative example of this challenge. Police officers are given the option to assign structured contributing-factor codes to each crash, while accompanying narratives describe the event in natural language. Discrepancies frequently arise due to ambiguity, reporting practices, and limitations of predefined coding schemes. As a result, structured labels may introduce systematic biases, omit relevant information, or encode interpretations not replicated in the narrative. Recent work has established that human disagreement in annotation is pervasive across NLP tasks (Pavlick and Kwiatkowski, 2019; Uma et al., 2021; Weerasooriya, 2024), and that quality management practices remain inconsistent even in major dataset creation efforts (Klie et al., 2024). At the same time, traditional multi-annotator approaches are often infeasible in restricted-access domains, where sensitive data cannot be widely shared and annotation requires specialized knowledge.

In this work, we propose a human-in-the-loop annotation framework designed to audit and analyze discrepancies between structured administrative labels and narrative-derived interpretations. Our approach combines (i) structured codes assigned during reporting, (ii) narrative-based annotation performed independently from structured codes, and (iii) expert adjudication guided by explicit uncertainty marking. This design enables controlled comparison across annotation layers while preserving data security. We evaluate this framework on a dataset of 9,387 crash narratives and conduct a comprehensive analysis of annotation discrepancies. Our study is guided by the following research questions:

- RQ1: How do discrepancies between structured administrative labels, narrative-based annotations, and adjudicated corrections reveal systematic annotation error and linguistic ambiguity in crash reporting?
- RQ2: What role does explicit uncertainty marking and human adjudication play in resolving annotation ambiguity and improving label consistency in high-stakes narrative corpora?

To address these questions, we perform five complementary analyses examining agreement patterns, discrepancy types, linguistic sources of ambiguity, and the role of uncertainty in annotation.

Our analysis reveals that 49.4% of records contain discrepancies between structured and narrative labels, with Distracted Driving exhibiting an 86.8% unsupported rate in structured codes. Narrative-based annotation achieves near-perfect agreement with adjudication ( $\kappa = 0.990$ ), and annotator-reported uncertainty predicts structured-code disagreement with an odds ratio of 8.91 ( $p < 10^{-124}$ ).

This work makes the following contributions:

- A human-in-the-loop annotation framework for restricted-access domains, combining single-annotator labeling, explicit uncertainty marking, and expert adjudication to ensure annotation quality under confidentiality constraints.
- A large-scale empirical analysis of annotation discrepancies across structured and narrative representations.
- A taxonomy of annotation discrepancies, distinguishing errors, omissions, ambiguity, and compression, and revealing class-specific patterns of annotation failure.
- An uncertainty-driven analysis of annotation difficulty, demonstrating that annotator-reported uncertainty strongly predicts both adjudication corrections and disagreement with structured codes.

## 2 Related Work

**Annotation Error and Quality.** Annotation quality directly shapes the reliability of NLP systems, yet systematic errors persist even in widely-used benchmarks. Northcutt et al. (2021) found an average of 3.4% label errors across 10 popular test

sets, demonstrating that label noise is pervasive rather than exceptional. Klie et al. (2023) reimplemented 18 annotation error detection methods and evaluated them across 9 datasets, establishing that no single method reliably detects errors across tasks. Weber-Genzel et al. (2024) introduced VariErr NLI, a two-round protocol that formally separates annotation errors from legitimate human label variation. Swayamdipta et al. (2020) proposed Dataset Cartography, showing that training dynamics can identify ambiguous and mislabeled instances. These works focus on annotator-produced labels in controlled settings. Our work extends this line of research to administratively produced labels created under operational constraints.

**Annotator Disagreement and Perspectivism.** A parallel line of research argues that disagreement between annotators carries meaningful signal rather than noise. Plank (2022) articulated this position comprehensively, arguing that human label variation impacts all stages of the ML pipeline and that the assumption of a single ground truth is often inappropriate. This view has been formalized as data perspectivism (Basile et al., 2021). Leonardelli et al. (2023) operationalized this view through the LeWiDi shared task at SemEval-2023, promoting models that learn from disagreement rather than resolving it through majority vote. Davani et al. (2022) showed that multi-task models predicting individual annotator ratings outperform majority-vote approaches on subjective tasks. Frenda et al. (2024) surveyed perspectivist approaches to NLP, documenting growing adoption of non-aggregated annotation across subjective tasks. However, existing disagreement research focuses almost exclusively on annotator-vs-annotator comparisons. Our work introduces a different type of comparison—annotation system vs. annotation system.

**Crash Narrative NLP.** NLP has been increasingly applied to crash narratives for automated safety analysis. Jaradat et al. (2024) demonstrated that text mining can uncover contributing factors from crash reports, while Bhagat et al. (2025) found that even expert-aligned LLM evaluation reveals persistent divergence between model predictions and human judgment on crash narratives. Oliace et al. (2023) applied BERT to injury classification from police reports. Despite this progress, all prior crash NLP work treats structured police codes as ground truth labels. No study has audited the annotation quality of these codes by comparing them

against independent narrative-derived labels—the gap our study addresses.

### 3 Methodology

We propose a multi-layer, human-in-the-loop annotation framework designed to systematically audit discrepancies between structured administrative labels and narrative-derived interpretations in high-stakes crash reporting. The overall workflow is illustrated in Figure 1.

#### 3.1 Problem Formulation

Let the dataset be defined as

$$\mathcal{D} = \{(x_i, s_i)\}_{i=1}^N \quad (1)$$

where  $x_i$  denotes a crash narrative and  $s_i \in \{0, 1\}^K$  represents the structured contributing-factor labels assigned by reporting officers, with  $K$  contributing-factor categories.

Our objective is not to predict labels, but to analyze annotation consistency across three layers:

- $s_i$ : structured labels (Layer 1)
- $n_i$ : narrative-based annotations (Layer 2)
- $a_i$ : adjudicated labels (Layer 3)

Each layer is represented as a multi-label vector:

$$s_i, n_i, a_i \in \{0, 1\}^K \quad (2)$$

To quantify differences between layers, we define a set-level discrepancy function:

$$\Delta(s_i, a_i) = \{(s_i \setminus a_i), (a_i \setminus s_i)\} \quad (3)$$

where  $s_i \setminus a_i$  denotes labels present in the structured set but absent from the adjudicated set (unsupported labels), and  $a_i \setminus s_i$  denotes labels present in the adjudicated set but absent from the structured set (omissions).

#### 3.2 Three-Layer Annotation Framework

To systematically study annotation discrepancies, we design a three-layer annotation process:

- **Layer 1 (Structured Labels):** Administrative contributing-factor codes selected by reporting officers as part of the crash report. These codes are not extracted from the narrative; they are recorded as separate structured fields in the reporting system.

- **Layer 2 (Narrative-Based Annotation):** Labels assigned independently using only officer written narrative text, with structured labels hidden to prevent bias.
- **Layer 3 (Adjudicated Labels):** Final labels obtained through expert-guided adjudication, representing an independent interpretation of narrative evidence.

This layered design enables controlled comparison between administrative reporting, text-derived interpretation, and expert validation, allowing us to isolate discrepancies arising from omission, ambiguity, and reinterpretation.

#### 3.3 Annotation Protocol

Annotation was conducted over approximately 14 months by a trained annotator, with quality maintained through regular review sessions with domain specialists from the City of Austin Transportation and Public Works department. In this context, domain specialists are the co-authors of this paper, who are transportation safety professionals from Austin Transportation and Public Works with practical experience reviewing crash reports, interpreting contributing-factor codes, and using crash data for traffic safety analysis. During these sessions, annotation decisions were audited and guidelines were iteratively refined to ensure consistency and domain alignment. Unlike conventional multi-annotator setups, we adopt a single-annotator-with-expert-oversight paradigm, motivated by:

- (a) data confidentiality constraints, (b) restricted access to sensitive records, and (c) the need for domain-specific expertise.

The annotation process follows three principles:

- **Narrative-Only Labeling:** Labels are assigned based solely on textual evidence.
- **Multi-Label Assignment:** Multiple contributing factors may be assigned to each narrative.
- **Explicit Uncertainty Marking:** Annotator flag cases where the narrative is ambiguous or insufficient for confident labeling.

Layer 2 annotation used only the officer-written narrative text; structured contributing-factor codes and other report fields were hidden to prevent label leakage. To support this workflow, we developed a custom annotation interface that presents narrative text to the annotator with structured codes hidden,

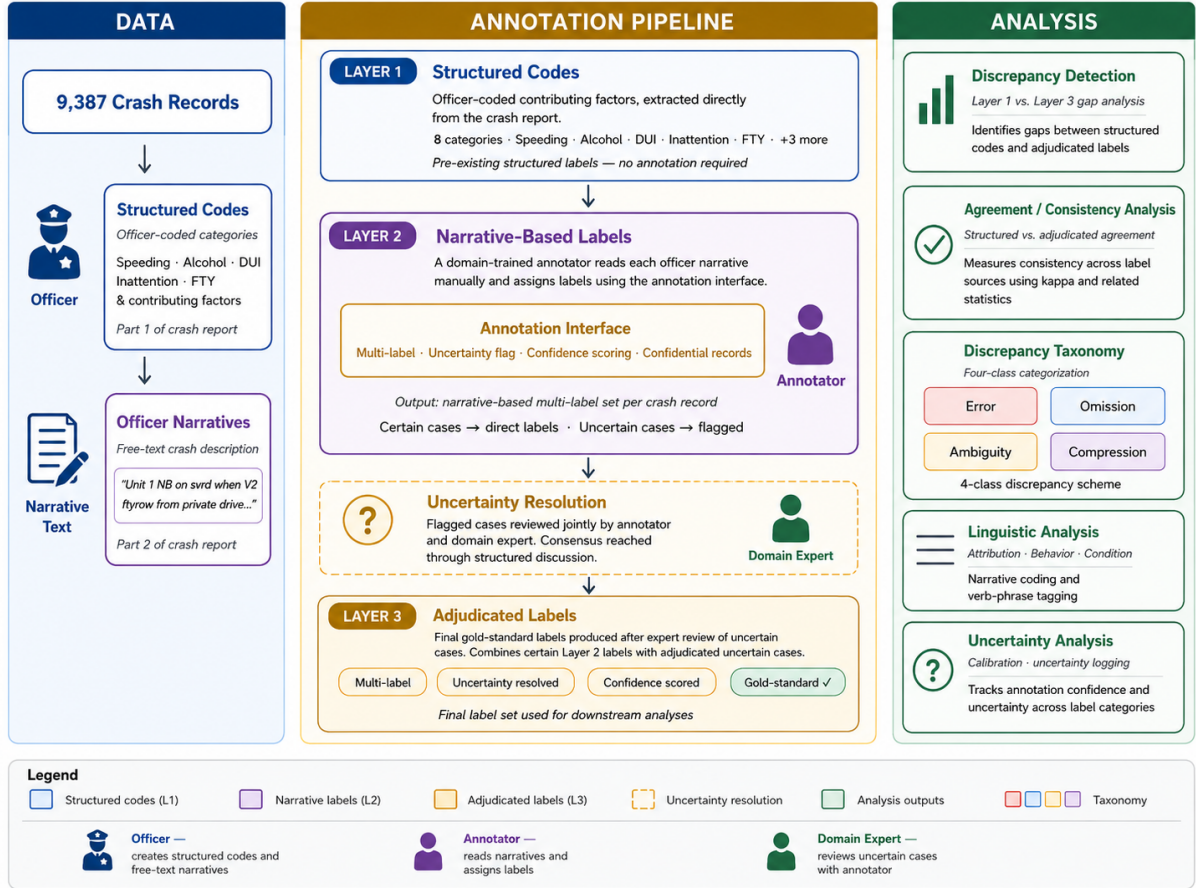


Figure 1: Overview of the proposed three-layer human-in-the-loop annotation framework. Structured officer labels (Layer 1) are contrasted with narrative-only annotations (Layer 2), while uncertain cases undergo expert adjudication to produce Layer 3. The resulting layers support discrepancy, taxonomy, linguistic, and uncertainty analyses.

supports multi-label selection, and includes an explicit uncertainty flag to capture ambiguous cases. A description of the interface and its design is provided in Appendix A.7.

### 3.4 Uncertainty-Guided Adjudication

Narratives marked as uncertain are routed to an adjudication stage, where the annotator and a domain expert jointly review the case. During adjudication, annotator has access to both the narrative and structured labels, allowing discrepancies to be examined and resolved in context.

Let  $u_i \in \{0, 1\}$  denote the uncertainty flag for sample  $i$ , where  $u_i = 1$  indicates ambiguity.

Adjudication is applied only to cases where the annotator explicitly marks uncertainty ( $u_i = 1$ ). During adjudication, labels may be:

(i) confirmed, (ii) revised, or (iii) left unresolved if insufficient evidence exists.

This process produces the adjudicated label vector  $a_i$ , which serves as a reference for evaluating annotation consistency.

The overall annotation workflow can be summarized as:

$$x_i \rightarrow n_i, u_i = \phi(x_i, n_i), \quad a_i = \begin{cases} \psi(x_i, n_i, s_i), & \text{if } u_i = 1 \\ n_i, & \text{otherwise} \end{cases} \quad (4)$$

### 3.5 Analytical Framework

To address our research questions, we design five complementary studies:

- Study 1: Comparison between structured labels  $s_i$  and narrative-based annotations  $n_i$
- Study 2: Comparison between narrative annotations  $n_i$  and adjudicated labels  $a_i$
- Study 3: Taxonomic classification of discrepancies  $\Delta(s_i, a_i)$
- Study 4: Linguistic analysis of ambiguity in narratives
- Study 5: Analysis of uncertainty signals  $u_i$  and their relationship to annotation correction

Each study provides a distinct perspective on annotation quality, discrepancy patterns, and ambiguity in high-stakes narrative data.

### 3.6 Label Space

We define a multi-label space with  $K = 8$  contributing-factor categories derived from an original set of 72 fine-grained police contributing-factor codes. The mapping was developed with domain specialists. This consolidation reduces sparsity while preserving interpretability. The categories include Speeding, Impaired Driving, Distracted Driving, Failure to Yield, Red Light Running, Access-Related Crashes, Visual Obstruction, and Other. Because both structured labels and narrative-based annotations were converted to the same 8-class label space before comparison, the consolidation itself did not create cross-layer disagreements. Rather, discrepancies reflect differences between the officer-selected structured codes and the narrative-derived labels within the same shared category scheme.

Detailed definitions, annotation guidelines, and edge cases for each category are provided in the Appendix A.6.

### 3.7 Data

We evaluate our approach on a real world dataset of police crash narratives obtained from the City of Austin Transportation and Public Works Department (TPW). The dataset contains 9,387 records. Each record contains two annotation sources: (i) a structured contributing-factor field with up to five codes selected from 72 predefined options, and (ii) a free-text narrative describing the crash circumstances in natural language.

## 4 Results and Analysis

### 4.1 Study 1: Structured vs. Narrative Agreement

To address RQ1, we compare structured contributing-factor labels with narrative-based annotations at the set level.

Across all records, 50.6% (4,747) show full agreement, while 49.4% (4,640) exhibit at least one discrepancy, indicating that structured labels fail to capture narrative evidence in nearly half of all cases. As shown in Figure 2, omission (present in narrative but absent in structured labels) and unsupported labeling (present in structured labels

without narrative evidence) patterns occur in 33.3% and 45.4% of records, respectively.

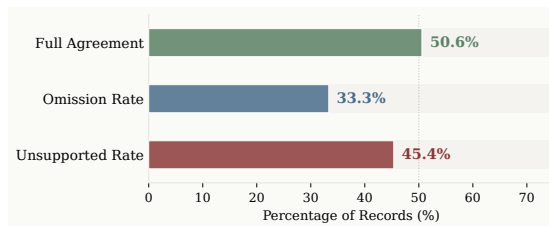


Figure 2: Overall structured vs. narrative agreement rates.

Agreement varies substantially across categories. As illustrated in Figure 3, the macro-average Cohen’s  $\kappa$  is 0.446, indicating moderate overall agreement. Clearly defined behaviors such as Red Light Running ( $\kappa = 0.696$ ), Failure to Yield ( $\kappa = 0.677$ ), and Speeding ( $\kappa = 0.624$ ) exhibit substantial agreement, while Distracted Driving ( $\kappa = 0.167$ ) and Access-Related Crashes ( $\kappa = 0.000$ ) show minimal to no agreement.

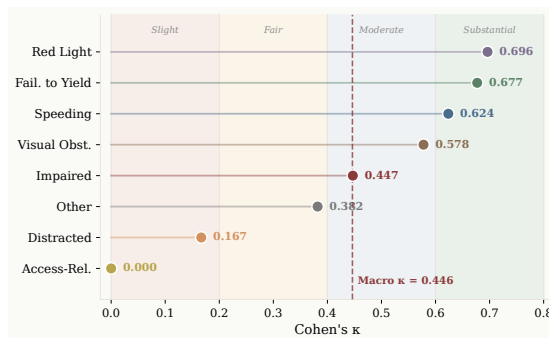


Figure 3: Per-class Cohen’s  $\kappa$  with interpretation bands.

To further characterize disagreement, we distinguish between omission and unsupported labeling. Figure 4 shows that these patterns are highly asymmetric across categories.

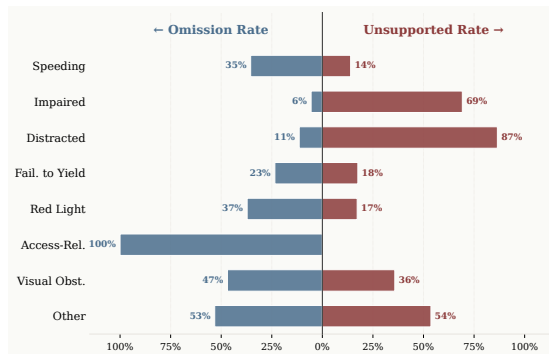


Figure 4: Per-class omission and unsupported rates.

In particular, Distracted Driving (87% unsp-

ported) and Impaired Driving (69% unsupported) are frequently assigned without supporting narrative evidence, while Access-Related Crashes (100% omission) and Red Light Running (37% omission) are often present in narratives but absent in structured codes.

A detailed confusion analysis in Appendix A.1 further reveals that certain structured labels consistently replace alternative contributing factors, indicating systematic substitution rather than random noise.

## 4.2 Study 2: Narrative vs. Adjudicated Agreement

To address RQ2, we compare narrative-based annotations with adjudicated labels.

Of all records, 852 (9.1%) were marked uncertain, while 8,535 (90.9%) were annotated with confidence.

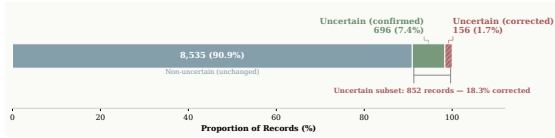


Figure 5: Distribution of annotation outcomes (uncertain vs. non-uncertain).

As shown in Figure 5, overall agreement between narrative and adjudicated labels is 98.3%, corresponding to a 1.7% correction rate across the full dataset.

This distribution indicates that uncertainty marking is conservative: The annotator flags potential ambiguity early, and most decisions remain valid under expert review. The observed corrections therefore reflect genuine annotation difficulty rather than systematic error.

Agreement remains consistently high across all categories. As reported in Table 1, the macro-average Cohen’s  $\kappa$  is 0.990, indicating near-perfect agreement. All classes achieve  $\kappa \geq 0.966$ , with Visual Obstruction (1.000) showing perfect agreement and Other (0.966) exhibiting the lowest consistency.

The contrast with Study 1 ( $\kappa = 0.446$ ) is substantial, indicating that disagreement primarily arises from structured administrative labels rather than inconsistencies in narrative-based interpretation. Adjudication produces targeted and asymmetric corrections. As illustrated in Figure 6, the largest net increase occurs for Failure to Yield

Category	$\kappa$
Visual Obstruction	1.000
Distracted Driving	0.998
Impaired Driving	0.997
Red Light Running	0.997
Access-Related Crashes	0.991
Speeding	0.989
Failure to Yield	0.983
Other	0.966
Macro Average	0.990

Table 1: Per-class agreement between narrative and adjudicated labels.

(+42), followed by Other (+24) and Speeding (+14), while Red Light Running shows a net decrease (-4).

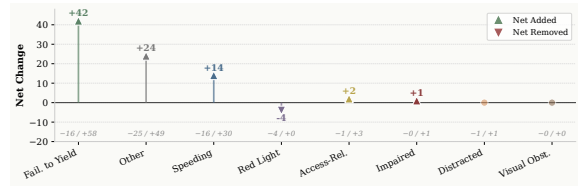


Figure 6: Net label changes during adjudication (additions vs. removals).

A detailed analysis of correction pathways (Appendix A.2) shows that adjudication primarily resolves ambiguous labels into more specific interpretations, with Other  $\rightarrow$  Failure to Yield being the most frequent transition.

## 4.3 Study 3: Discrepancy Taxonomy

To move beyond aggregate agreement metrics, we classify discrepancies between structured codes and adjudicated labels into a four-category taxonomy aligned with the LAW special theme of errors in annotation: Error (unsupported structured label), Omission (missing structured label), Ambiguity (insufficient narrative evidence), and Compression (structurally unrepresentable cases). Categories are not mutually exclusive.

In this study, an unsupported label refers to a structured contributing-factor code that is present in Layer 1 but not supported by the narrative evidence after Layer 3 adjudication. This does not necessarily imply that the officer made a factual mistake; rather, it indicates that the structured code could not be verified from the narrative text under our annotation guidelines. Our analysis therefore treats the narrative as the evidence source for narrative-based annotation, not as a complete reconstruction of the officer’s knowledge or intent during reporting.

Of all records, 50.6% (4,754) show no discrep-

Category	Records	% of All	% of Discrepancies
None	4,754	50.6%	—
Error	4,225	45.0%	91.2%
Omission	2,911	31.0%	62.8%
Ambiguity	799	8.5%	17.2%
Compression	238	2.5%	5.1%

Table 2: Discrepancy taxonomy distribution. Categories can overlap; percentages of discrepancies sum to more than 100%.

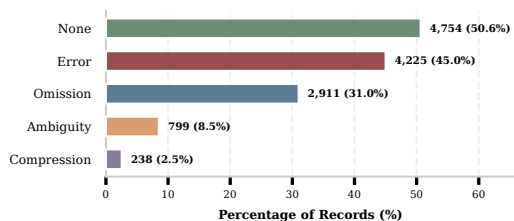


Figure 7: Distribution of discrepancy taxonomy categories.

ancy, while 49.4% (4,633) exhibit at least one discrepancy. As shown in Figure 7 and Table 2, Error is the dominant discrepancy type, appearing in 45.0% of all records and 91.2% of discrepant cases, followed by Omission (31.0%). Ambiguity (8.5%) and Compression (2.5%) occur less frequently but reflect qualitatively different failure modes.

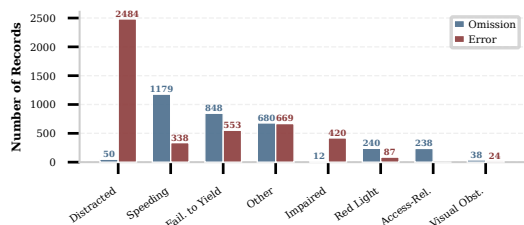


Figure 8: Per-class breakdown of error vs. omission discrepancies.

As shown in Figure 8, Distracted Driving is dominated by errors (2,484) with minimal omissions (50), confirming systematic over-coding in structured labels. Speeding and Failure to Yield show more balanced profiles, while Access-Related Crashes are entirely attributable to omission, reflecting a structural limitation of the coding system.

A consolidated view of annotation fragility is shown in Figure 9, which ranks classes by discrepancy rate. These results show that discrepancies are not uniformly distributed but concentrated in specific classes and driven by distinct mechanisms. In particular, structured over-coding is the dominant source of disagreement, indicating that annotation

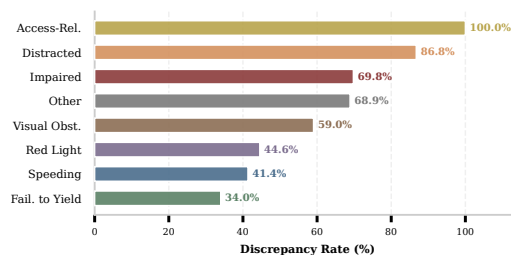


Figure 9: Per-class discrepancy rates across contributing factors.

inconsistencies arise primarily from unsupported structured labels rather than missing information.

A detailed per-class breakdown of discrepancy behavior across classes is provided in Appendix A.3.

#### 4.4 Study 4: Linguistic Sources of Ambiguity

To understand why discrepancies arise, we identify five recurring linguistic phenomena that systematically introduce annotation ambiguity: hedging, implicit causality, underspecification, multi-event compression, and perspective markers. Here are some narrative examples.<sup>1</sup>

*“Unit 1 stopped at a stop sign before entering the intersection. Unit 2 proceeded through on the ABC street. Both drivers indicated a stop sign lying on the ground that would have controlled Unit 2’s lane.”*

Here, the downed sign’s causal role is never confirmed, the narrative does not establish whether it fell before or after the collision (implicit causality), and the key evidence is attributed to the drivers rather than verified by the officer (perspective markers), leaving the annotator to choose between Failure to Yield and Other based on ambiguous evidence.

*“Unit 1 was blocked in a parking stall. The driver possibly backed up slightly and may have collided with Unit 2, then went forward in a panic and struck Unit 3.”*

Here, hedging expressions (“possibly,” “may have”) signal that even the reporting officer could not confirm whether the first collision occurred, forcing the annotator to assign labels based on unverified events. Additionally, the narrative describes two separate collisions in sequence, but the

<sup>1</sup>Due to the data privacy agreement, all police narrative examples referenced in this paper are paraphrased and de-identified.

Metric	Uncertain	Non-uncertain
Correction Rate	18.3%	0.0%
Discrepancy Rate	88.1%	45.5%

Table 3: Correction rate and structured-adjudicated discrepancy rate by uncertainty status.

	Discrepancy	Agreement	Total
Uncertain	751	101	852
Non-uncertain	3,882	4,653	8,535
Total	4,633	4,754	9,387

Table 4: Cross-tabulation of uncertainty status and structured-adjudicated discrepancy.

structured coding system captures only a single contributing factor (multi-event compression).

*“Unit 1 was traveling southbound in the far left lane when it struck the left side cable barrier.”*

Here, the narrative provides no explanation for why the vehicle departed its lane, the annotator cannot determine whether the cause was Speeding, Distracted Driving, Impaired Driving, or a mechanical failure, as the description lacks any behavioral or contextual detail (underspecification).

These phenomena are not random sources of noise but reflect systematic properties of narrative reporting based on incomplete information. Detailed examples and analysis for all five phenomena are provided in Appendix A.4.

#### 4.5 Study 5: Uncertainty Analysis

We evaluate whether annotator-reported uncertainty reflects meaningful annotation difficulty. Among all records, 852 (9.1%) were marked uncertain, while 8,535 (90.9%) were annotated with confidence. As shown in Table 3, all observed corrections occur within the uncertain subset because adjudication was applied only to records explicitly marked as uncertain. Non-uncertain records were carried forward unchanged by design. Uncertain records also exhibit substantially higher disagreement with structured codes (88.1% vs. 45.5%).

This relationship is statistically significant. As summarized in Table 4, a chi-square test yields  $\chi^2 = 562.37$  ( $p < 10^{-124}$ ), with an odds ratio of 8.91, indicating that uncertain records are nearly nine times more likely to exhibit disagreement with structured labels.

Uncertainty is unevenly distributed across classes. As shown in Figure 10, the Other category exhibits the highest uncertainty rate (30.6%), followed by Failure to Yield (10.0%) and Visual Obstruction (9.9%). In contrast, Distracted Driv-

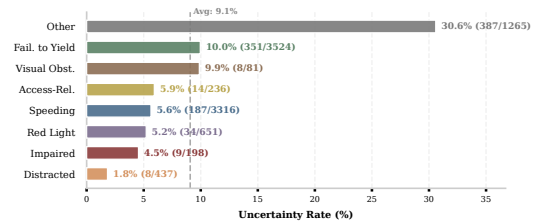


Figure 10: Per-class concentration of uncertainty in annotation.

ing (1.8%) and Impaired Driving (4.5%) show low uncertainty despite high structured-code error rates (Study 3), indicating that these factors are linguistically explicit in narratives but systematically underrepresented in structured coding.

## 5 Discussion

Our results show that annotation discrepancies in crash narratives are systematic rather than random. Importantly, these discrepancies should be interpreted as differences between administrative coding and narrative-grounded annotation, rather than as direct evidence of officer error. Structured crash codes may reflect reporting conventions, agency requirements, or contextual information available to the officer but not explicitly stated in the narrative. Label categories also differ in how directly they can be inferred from narrative text. For example, impaired driving may be supported by explicit evidence such as alcohol or drug involvement, whereas distracted driving often depends on more indirect behavioral descriptions such as inattention, failure to perceive hazards, or phone use. These differences help explain why some categories exhibit higher ambiguity or unsupported-label rates than others. Most disagreements are driven by unsupported structured labels, suggesting that administrative coding practices can introduce bias rather than simply omit information. At the same time, narrative-based annotation combined with uncertainty-guided adjudication achieves near-perfect agreement ( $\kappa = 0.990$ ), demonstrating that reliable annotation is possible even with a single trained annotator when decisions are grounded in textual evidence and expert review. We further show that many disagreements arise from linguistic ambiguity, such as hedging and underspecification, highlighting limitations of purely rule-based annotation. Importantly, annotator-reported uncertainty emerges as a strong signal of annotation dif-

ficuity, effectively identifying cases that require expert adjudication. Together, these findings support a human-in-the-loop annotation paradigm for high-stakes, restricted-access domains. This framework generalizes to other high-stakes settings, such as clinical records, legal documents, and safety-critical reporting systems, where data access and domain expertise impose similar constraints.

## 6 Conclusion and Future Work

We present a human-in-the-loop framework for analyzing annotation discrepancies in crash narratives using structured labels, narrative annotation, and expert adjudication. Our results show that discrepancies are systematic and largely driven by unsupported structured labels and linguistic ambiguity. Across five studies, we demonstrate that narrative-based annotation achieves high consistency under expert review, that discrepancies can be organized through a clear taxonomy, and that annotator-reported uncertainty is a reliable signal of annotation difficulty. These findings highlight limitations of administrative coding and the importance of text-grounded, human-centered annotation. More broadly, this work offers a practical annotation paradigm for restricted-access domains, where traditional multi-annotator approaches are not feasible.

For future work, we plan to extend this framework in several directions. First, we aim to explore semi-automated annotation support, including models that predict uncertainty and assist annotators during labeling. Second, we will investigate error-aware training strategies that leverage discrepancy patterns to improve downstream predictive models. Third, subject to data-sharing constraints, we intend to release a curated subset of annotated data and the annotation tool to support reproducibility and further research.

## 7 Limitations

This study also has several limitations. First, annotation was carried out by a single trained annotator due to restricted data access and the need for domain expertise. Although this design was motivated by confidentiality constraints and the need for domain expertise, it may still introduce annotator-specific bias. To partially assess consistency, the annotator re-labeled a random subset of records after a time gap, and the results were reviewed for stability. Future work should extend this check to

a larger systematically sampled subset and report formal intra-annotator agreement statistics. Instead of conventional multi-annotator redundancy, our framework relies on uncertainty-guided adjudication and expert verification. Second, the dataset is confidential and cannot be fully shared, which limits reproducibility, although we mitigate this by providing detailed guidelines and taxonomy definitions. Third, the taxonomy compresses 72 original contributing-factor codes into 8 classes, which may reduce some fine-grained distinctions. Finally, adjudication was applied only to uncertain cases, reflecting a practical workflow but potentially underestimating total disagreement.

## Ethics Statement

This study uses 9,387 crash records provided by a municipal transportation agency under a formal data-sharing agreement for traffic safety research. As the data is administrative and not directly collected from human subjects, IRB approval was not required. All work complies with the ACL Code of Ethics. To protect privacy, no personally identifiable information was included in the dataset, all processing was conducted on secure local infrastructure, and no raw data was shared with external services or commercial models. All examples are paraphrased to prevent identification of real incidents. The system is designed for safety analysis, and while dual-use risks exist, its primary goal is to improve data quality for infrastructure-level decision-making.

## Acknowledgments

The authors acknowledge the support provided by the U.S. Department of Transportation through the National Center for Infrastructure Transformation under grant numbers 69A3552344813 and 69A3552348318. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the U.S. Government.

## References

Valerio Basile, Michael Fell, Tommaso Fornaciari, Dirk Hovy, Silviu Paun, Barbara Plank, Massimo Poesio, and Alexandra Uma. 2021. We need to consider disagreement in evaluation. In *Proceedings of the 1st Workshop on Benchmarking: Past, Present and Future*, pages 15–21.

- Sudesh Ramesh Bhagat, Ibne Farabi Shihab, and Anuj Sharma. 2025. Accuracy is not agreement: Expert-aligned evaluation of crash narrative classification models. In *arXiv preprint arXiv:2504.13068*.
- Sharon Campbell and Katrina Giadresco. 2020. Computer-assisted clinical coding: A narrative review of the literature on its benefits, limitations, implementation and impact on clinical coding professionals. *Health Information Management Journal*, 49(1):5–18.
- Aida Mostafazadeh Davani, Mark Díaz, and Vinodkumar Prabhakaran. 2022. Dealing with disagreements: Looking beyond the majority vote in subjective annotations. *Transactions of the Association for Computational Linguistics*, 10:92–110.
- Beniamino Di Martino, Fiammetta Marulli, Pietro Lupi, and Alessandra Cataldi. 2020. A machine learning based methodology for automatic annotation and anonymisation of privacy-related items in textual documents for justice domain. In *Conference on Complex, Intelligent, and Software Intensive Systems*, pages 530–539. Springer.
- Simona Frenda, Gavin Abercrombie, Valerio Basile, Alessandro Pedrani, Raffaella Panizzon, Alessandra Teresa Cignarella, Cristina Marco, and Davide Bernardi. 2024. Perspectivist approaches to natural language processing: A survey. *Language Resources and Evaluation*.
- Shadi Jaradat, Taqwa I. Alhadidi, Huthaifa I. Ashqar, Ahmed Hossain, and Mohammed Elhenawy. 2024. Exploring traffic crash narratives in Jordan using text mining analytics. *2024 IEEE 3rd International Conference on Computing and Machine Intelligence (ICMI)*, pages 1–6.
- Wirdatul Khasanah, Hale Yılmaz, and Benjamin White. 2025. The validity of automated essay scoring using NLP compared to human raters in the context of language certification exams. *JILTECH: Journal International of Lingua & Technology*, 4(3).
- Jan-Christoph Klie, Richard Eckart de Castilho, and Iryna Gurevych. 2024. Analyzing dataset annotation quality management in the wild. *Computational Linguistics*, 50(3):817–866.
- Jan-Christoph Klie, Bonnie Webber, and Iryna Gurevych. 2023. Annotation error detection: Analyzing the past and present for a more coherent future. *Computational Linguistics*, 49(1):157–198.
- Akshi Kumar and Saurabh Raj Sangwan. 2025. Introduction to natural language processing in high-stakes domains. In *Transformative Natural Language Processing: Bridging Ambiguity in Healthcare, Legal, and Financial Applications*, pages 1–22. Springer.
- Elisa Leonardelli, Gavin Abercrombie, Dina Almanea, Valerio Basile, Tommaso Fornaciari, Barbara Plank, Verena Rieser, Alexandra Uma, and Massimo Poesio. 2023. Semeval-2023 task 11: Learning with disagreements (LeWiDi). In *Proceedings of the 17th International Workshop on Semantic Evaluation (SemEval-2023)*, pages 2304–2318, Toronto, Canada.
- Curtis G. Northcutt, Anish Athalye, and Jonas Mueller. 2021. Pervasive label errors in test sets destabilize machine learning benchmarks. In *Proceedings of the 35th Conference on Neural Information Processing Systems (NeurIPS), Datasets and Benchmarks Track*.
- Amir Hossein Oliaee, Subasish Das, Jinli Liu, and M. Ashifur Rahman. 2023. Using bidirectional encoder representations from transformers (BERT) to classify traffic crash severity types. *Natural Language Processing Journal*, 3(S1).
- Ellie Pavlick and Tom Kwiatkowski. 2019. Inherent disagreements in human textual inferences. *Transactions of the Association for Computational Linguistics*, 7:677–694.
- Barbara Plank. 2022. The “problem” of human label variation: On ground truth in data, modeling and evaluation. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 10671–10682, Abu Dhabi, United Arab Emirates.
- Aishat O. Salami. 2023. Leveraging natural language processing to detect non-compliance in clinical documentation: Current advances, challenges, and future directions. *International Journal of Scientific Research in Science, Engineering and Technology*, pages 459–473.
- Swabha Swayamdipta, Roy Schwartz, Nicholas Lourie, Yizhong Wang, Hannaneh Hajishirzi, Noah A. Smith, and Yejin Choi. 2020. Dataset cartography: Mapping and diagnosing datasets with training dynamics. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 9275–9293.
- Alexandra N. Uma, Tommaso Fornaciari, Dirk Hovy, Silviu Paun, Barbara Plank, and Massimo Poesio. 2021. Learning from disagreement: A survey. *Journal of Artificial Intelligence Research*, 72:1385–1470.
- Leon Weber-Genzel, Siyao Peng, Marie-Catherine de Marneffe, and Barbara Plank. 2024. VariErr NLI: Separating annotation error from human label variation. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 2256–2269, Bangkok, Thailand.
- Tharindu Cyril Weerasooriya. 2024. *Learning from Disagreement in Human-Annotated Datasets*. Ph.D. thesis, Rochester Institute of Technology.

## A Appendix

### A.1 Study 1: Extended Analysis of Structured vs. Narrative Discrepancies

This section provides additional analysis supporting the results presented in Section 4.1.

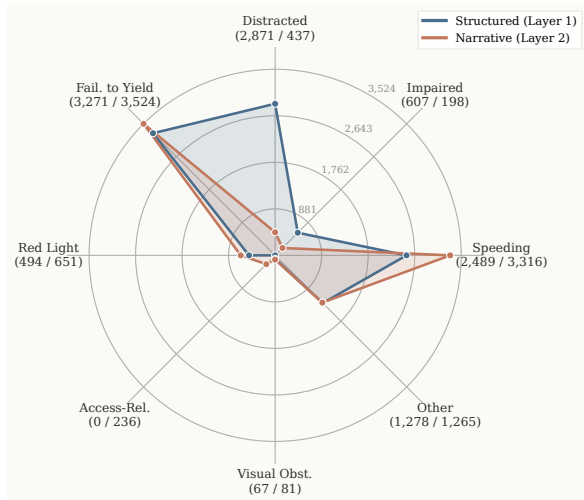


Figure 11: Label distribution comparison between structured and narrative annotations.

**Label Distribution Analysis.** Figure 11 compares the distribution of contributing-factor labels across structured and narrative annotations. The chart highlights substantial imbalances across several categories. In particular, Distracted Driving is heavily overrepresented in structured labels (2,871) compared to narrative annotations (437), while Access-Related Crashes are present in narratives (236) but entirely absent from structured coding. These patterns indicate systematic differences in how contributing factors are represented, with certain categories being over-assigned in structured data and others omitted entirely.

**Confusion Matrix Analysis.** Figure 12 provides a detailed view of label transitions between structured and narrative annotations. The matrix reveals systematic substitution patterns beyond simple disagreement. High agreement is observed for categories such as Speeding (91%) and Failure to Yield (88%), while Distracted Driving shows only 18% agreement. In cases of disagreement, Distracted Driving labels are frequently reassigned to other categories, most commonly:

- Speeding (32%)
- Failure to Yield (23%)

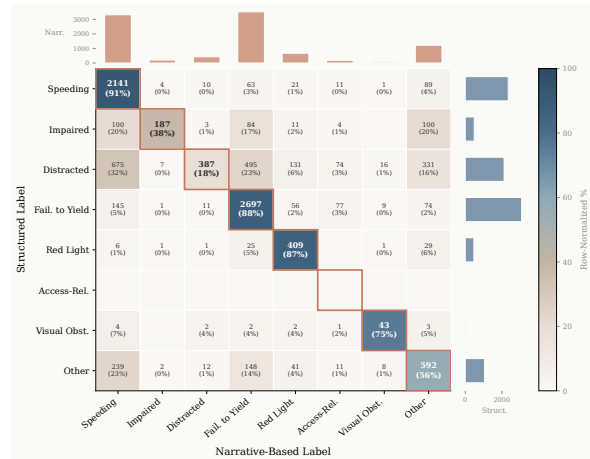


Figure 12: Row-normalized confusion matrix with marginal distributions: structured  $\rightarrow$  narrative labels.

- Other (16%)

These patterns indicate that certain structured labels do not merely introduce noise but systematically replace alternative contributing factors.

#### Additional Observations.

- Categories with low agreement (e.g., Distracted Driving) correspond to high unsupported labeling rates observed in the main analysis.
- Categories with high omission rates (e.g., Access-Related Crashes) are consistently underrepresented in structured data.

### A.2 Study 2: Extended Analysis of Adjudication Effects

This section provides additional analyses of label corrections and class-level transitions during adjudication, complementing the summary results in Section 4.2.

Figure 13 presents the correction flow matrix between narrative and adjudicated labels, revealing structured transition patterns during adjudication. The most frequent transition is from Other  $\rightarrow$  Failure to Yield (17 cases), indicating that ambiguous cases are often resolved into more specific contributing factors. The reverse transition (Failure to Yield  $\rightarrow$  Other, 7 cases) occurs less frequently, confirming that adjudication tends to refine rather than generalize labels.

### A.3 Study 3: Extended Taxonomy Analysis

This section provides additional detail on the discrepancy taxonomy introduced in Section 4.3, in-

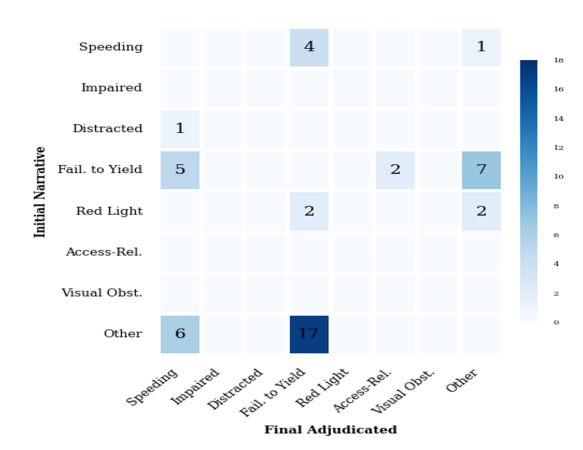


Figure 13: Correction flow between narrative and adjudicated labels. Each cell represents the number of transitions from an initial narrative label (rows) to a final adjudicated label (columns) within the uncertain subset.

cluding formal definitions and class-level breakdowns.

**Taxonomy Definitions.** Each discrepancy between structured and adjudicated labels is assigned one or more of the following categories:

- **Error:** A structured label is present but unsupported by narrative evidence.
- **Omission:** A contributing factor is present in the narrative but absent from structured codes.
- **Ambiguity:** The adjudicated label includes Other and a discrepancy exists, indicating insufficient narrative clarity.
- **Compression:** The narrative contains an information that cannot be represented in the structured coding system.

These categories are not mutually exclusive, and a single record may exhibit multiple discrepancy types.

**Detailed Per-Class Discrepancy Breakdown.** Table 5 provides a detailed view of discrepancy behavior across classes. The results confirm strong asymmetry in several categories:

- Distracted Driving is dominated by errors (2,484) with minimal omissions (50), indicating systematic over-coding in structured labels.
- Impaired Driving shows a similar pattern (420 errors vs. 12 omissions).

- Speeding and Failure to Yield exhibit more balanced discrepancies, with both omission and error contributing substantially.
- Other shows near-symmetric behavior (680 omissions vs. 669 errors), consistent with its role as a residual category.

#### Additional Observations.

- Error is the dominant discrepancy type across nearly all classes, reinforcing that structured labels frequently introduce unsupported information.
- Omission remains substantial in several categories, particularly Speeding and Failure to Yield, indicating that structured codes also fail to capture important narrative evidence.
- Compression is confined to Access-Related Crashes, reflecting a limitation of the reporting schema rather than annotator behavior.

#### A.4 Study 4: Extended Linguistic Ambiguity Analysis

This section provides detailed examples of linguistic phenomena that introduce annotation ambiguity, complementing the summary presented in Section 4.4.

**Hedging Language. Example:** “Unit 1 appeared to be traveling at an excessive speed before losing control.”

Hedging expressions (e.g., “appeared to be”) signal that the reported behavior is inferred rather than directly observed. This creates uncertainty in determining whether the evidence is sufficient to assign a label such as Speeding. In practice, structured labels often treat such inferences as definitive, while narrative-based annotation may apply stricter evidence thresholds.

**Implicit Causality. Example:** “The roadway was wet from recent rainfall. Unit 2 crossed the center line and struck Unit 1.”

Environmental conditions are described but not explicitly linked to causation. The annotator must decide whether such conditions constitute contributing factors or merely contextual information, leading to inconsistent interpretations.

Contributing Factor	Omitted by Police	Erroneously Coded	Total	Disc. Rate
Distracted Driving	50	2,484	2,534	86.8%
Other	680	669	1,349	68.9%
Impaired Driving	12	420	432	69.8%
Speeding	1,179	338	1,517	41.4%
Failure to Yield	848	553	1,401	34.0%
Red Light Running	240	87	327	44.6%
Access-Related Crashes	238	0	238	100.0%
Visual Obstruction	38	24	62	59.0%

Table 5: Per-class discrepancy breakdown showing omissions and errors for each contributing factor.

**Underspecification.** **Example:** “Unit 1 was traveling northbound on the highway when it drifted across the center line and struck Unit 2 head-on. Unit 1’s driver was transported to the hospital. No further information was available at the scene.” The narrative describes the crash outcome but provides no behavioral, environmental, or mechanical context for why the vehicle crossed the center line. The annotator cannot distinguish between Speeding, Distracted Driving, Impaired Driving, or a medical event, as the description omits the causal detail needed for confident classification. The phrase “no further information was available” confirms that the underspecification is inherent to the report itself, not an oversight by the annotator.

**Multi-Event Compression.** **Example:** “Unit 1 was using a cell phone, failed to notice stopped traffic, and struck Unit 2 before veering into another lane.”

Multiple contributing factors are described, but structured coding often captures only one. This results in systematic loss of information and contributes to discrepancies between structured and narrative annotations.

**Perspective Markers.** **Example:** “The witness stated that Unit 1 ran the red light, while the driver claimed the light was yellow.”

Conflicting accounts attributed to different sources introduce ambiguity. The annotator must determine which account to prioritize, reducing consistency and increasing uncertainty.

**Summary.** These phenomena, hedging, implicit causality, underspecification, multi-event compression, and perspective markers, represent recurring linguistic patterns that systematically produce annotation ambiguity. They reflect inherent properties of narrative reporting rather than annotation error alone.

## A.5 Study 5: Extended Uncertainty Analysis

This section provides detailed statistical and class-level analyses supporting the results in Section 4.5.

Uncertainty is unevenly distributed across contributing factor classes. The Other category exhibits the highest uncertainty rate (30.6%), followed by Failure to Yield (10.0%) and Visual Obstruction (9.9%), reflecting classes that require complex situational interpretation.

In contrast, Distracted Driving (1.8%) and Impaired Driving (4.5%) show low uncertainty despite high structured-code error rates (see Study 3). This indicates that these factors are linguistically explicit in narratives but systematically misrepresented in structured coding.

### Additional Observations.

- Uncertainty is concentrated in classes requiring contextual reasoning.
- Linguistically explicit behaviors exhibit low uncertainty even when structured labels are unreliable.
- This supports the use of uncertainty as a targeted signal for annotation difficulty.

## A.6 Annotation Guidelines and Label Definitions

We follow a three-layer annotation framework in which crash narratives are labeled using an 8-class contributing-factor taxonomy derived from an original set of 72 police codes. The mapping was developed in collaboration with domain specialists to reduce sparsity while preserving interpretability.

**General Annotation Principles.** Annotations are assigned based solely on textual evidence in the narrative, without reference to structured administrative labels. Multiple contributing factors

Contributing Factor	Total Records	Uncertain	Uncertainty Rate
Other	1,265	387	30.6%
Failure to Yield	3,524	351	10.0%
Visual Obstruction	81	8	9.9%
Access-Related Crashes	236	14	5.9%
Speeding	3,316	187	5.6%
Red Light Running	651	34	5.2%
Impaired Driving	198	9	4.5%
Distracted Driving	437	8	1.8%

Table 6: Per-class uncertainty concentration.

may be assigned when supported by the text. When evidence is insufficient or ambiguous, records are explicitly flagged using an uncertainty indicator. All examples used in this study are paraphrased to preserve data confidentiality.

**Label Definitions.** We summarize the eight contributing-factor classes below:

- **Speeding:** The driver operated the vehicle at an unsafe speed or followed too closely, including failure to control speed or rear-end collisions where no alternative cause is specified.
- **Impaired Driving:** The driver was under the influence of alcohol, drugs, or medication, with observable evidence linking impairment to the crash.
- **Distracted Driving:** The driver’s attention was diverted (e.g., phone use, inattention, fatigue), resulting in failure to perceive or react to hazards.
- **Failure to Yield:** The driver violated right-of-way rules or executed unsafe maneuvers such as improper turns, lane changes, or passing.
- **Red Light Running:** The driver entered an intersection in violation of a red signal or stop control.
- **Visual Obstruction:** The driver’s visibility was physically blocked by environmental or structural factors (e.g., glare, weather, obstacles).
- **Access-Related Crash:** The crash occurred while entering or exiting a private complex, parking lot, or private access point.

- **Other:** Residual category for contributing factors not captured by the above classes (e.g., animals, medical events, vehicle issues, road rage).

These definitions reflect a balance between semantic clarity and compatibility with existing reporting standards.

The original reporting system includes 72 fine-grained contributing-factor codes, which were consolidated into the 8-class taxonomy described above. The mapping from 72 original contributing-factor codes to the 8-class taxonomy was developed in collaboration with domain experts, ensuring alignment with real-world reporting practices. The complete code-to-class mapping is provided in Table 7.

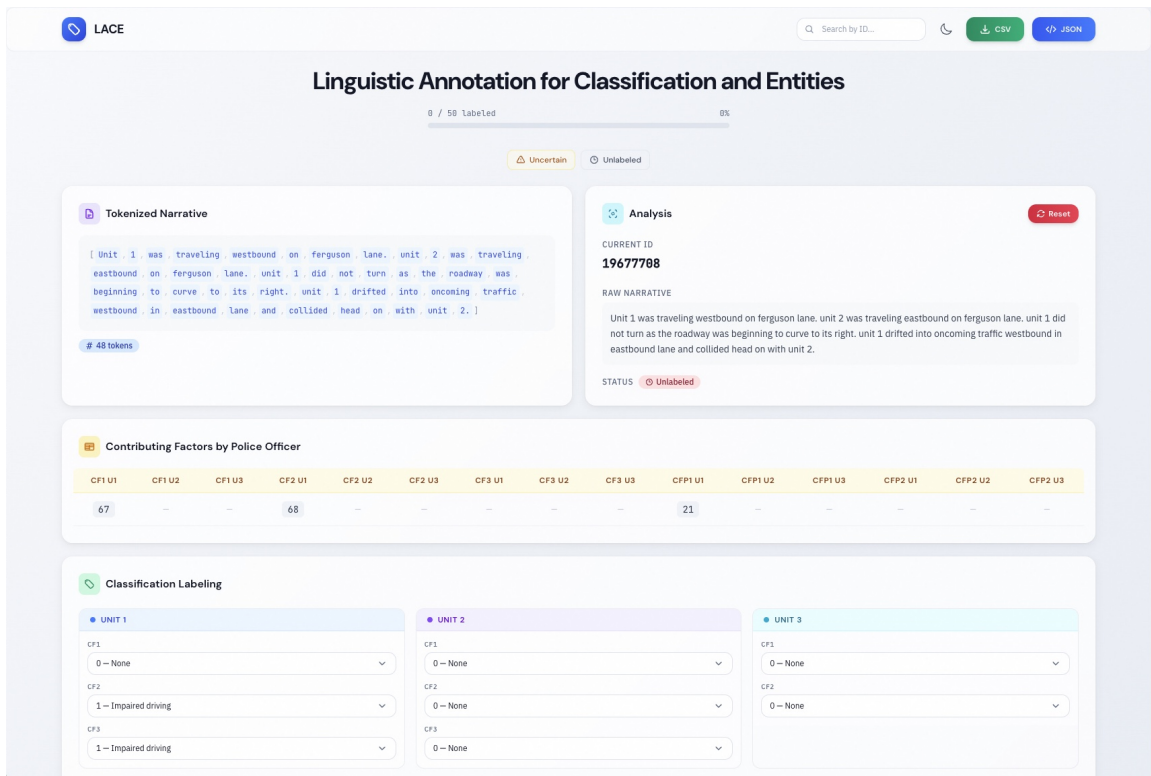
### A.7 Annotation Interface

We developed a custom annotation interface to support multi-label assignment and explicit uncertainty marking in crash narratives. The tool allows annotators to select contributing factors, flag ambiguous cases, and review previously labeled records within a unified workflow.

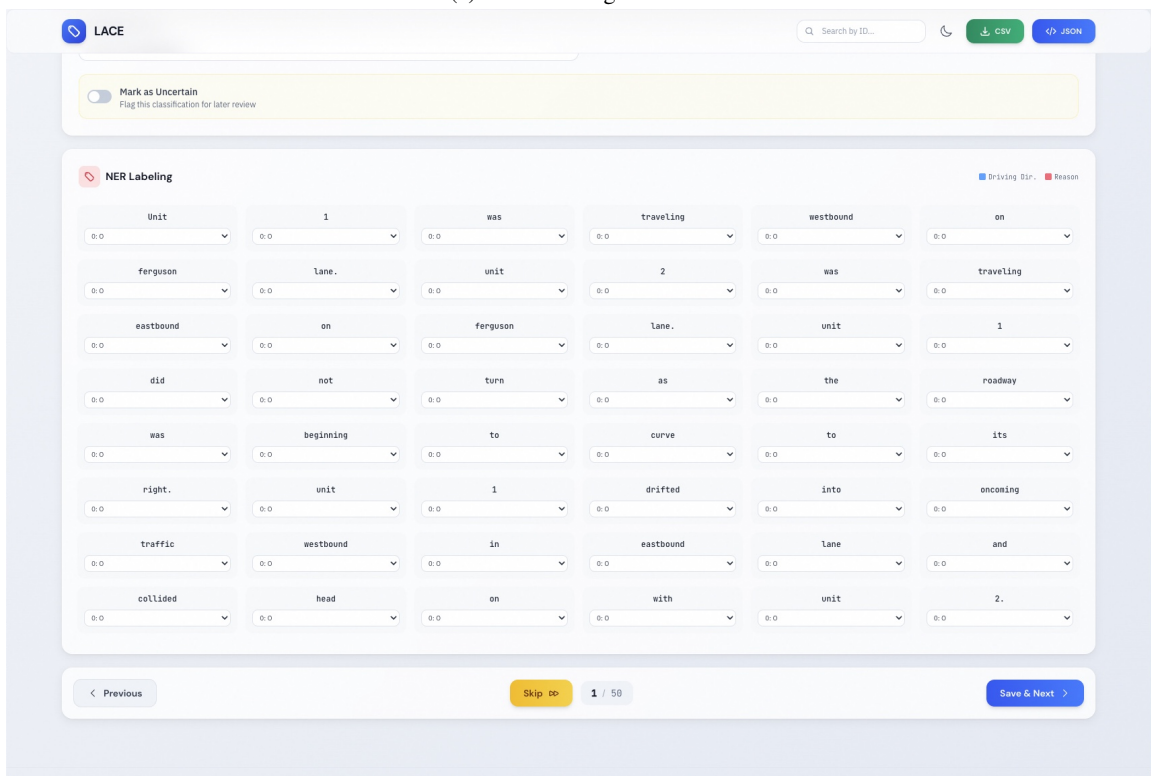
The interface was designed to operate in a restricted-access environment to ensure data confidentiality while maintaining annotation consistency.

ID	Original Contributing Factor Code	8-Class Category
0	NONE	—
1	ANIMAL ON ROAD - DOMESTIC	Other
2	ANIMAL ON ROAD - WILD	Other
3	BACKED WITHOUT SAFETY	Failure to Yield
4	CHANGED LANE WHEN UNSAFE	Failure to Yield
14	DISABLED IN TRAFFIC LANE	Other
15	DISREGARD STOP AND GO SIGNAL	Red Light Running
16	DISREGARD STOP SIGN OR LIGHT	Red Light Running
17	DISREGARD TURN MARKS AT INTERSECTION	Failure to Yield
18	DISREGARD WARNING SIGN AT CONSTRUCTION	Failure to Yield
19	DISTRACTION IN VEHICLE	Distracted Driving
20	DRIVER INATTENTION	Distracted Driving
21	DROVE WITHOUT HEADLIGHTS	Other
22	FAILED TO CONTROL SPEED	Speeding
23	FAILED TO DRIVE IN SINGLE LANE	Failure to Yield
24	FAILED TO GIVE HALF OF ROADWAY	Failure to Yield
25	FAILED TO HEED WARNING SIGN	Failure to Yield
26	FAILED TO PASS TO LEFT SAFELY	Failure to Yield
27	FAILED TO PASS TO RIGHT SAFELY	Failure to Yield
28	FAILED TO SIGNAL OR GAVE WRONG SIGNAL	Failure to Yield
29	FAILED TO STOP AT PROPER PLACE	Failure to Yield
30	FAILED TO STOP FOR SCHOOL BUS	Failure to Yield
31	FAILED TO STOP FOR TRAIN	Failure to Yield
32	FAILED TO YIELD ROW - EMERGENCY VEHICLE	Failure to Yield
33	FAILED TO YIELD ROW - OPEN INTERSECTION	Failure to Yield
34	FAILED TO YIELD ROW - PRIVATE DRIVE	Failure to Yield
35	FAILED TO YIELD ROW - STOP SIGN	Failure to Yield
36	FAILED TO YIELD ROW - TO PEDESTRIAN	Failure to Yield
37	FAILED TO YIELD ROW - TURNING LEFT	Failure to Yield
38	FAILED TO YIELD ROW - TURN ON RED	Failure to Yield
39	FAILED TO YIELD ROW - YIELD SIGN	Failure to Yield
40	FATIGUED OR ASLEEP	Distracted Driving
41	FAULTY EVASIVE ACTION	Failure to Yield
42	FIRE IN VEHICLE	Other
43	FLEEING OR EVADING POLICE	Other
44	FOLLOWED TOO CLOSELY	Speeding
45	HAD BEEN DRINKING	Impaired Driving
46	HANDICAPPED DRIVER	Other
47	ILL	Other
48	IMPAIRED VISIBILITY	Visual Obstruction
49	IMPROPER START FROM PARKED POSITION	Other
50	LOAD NOT SECURED	Other
51	OPENED DOOR INTO TRAFFIC LANE	Failure to Yield
52	OVERSIZED VEHICLE OR LOAD	Other
53	OVERTAKE AND PASS INSUFFICIENT CLEARANCE	Failure to Yield
54	PARKED AND FAILED TO SET BRAKES	Other
55	PARKED IN TRAFFIC LANE	Other
56	PARKED WITHOUT LIGHTS	Other
57	PASSED IN NO PASSING LANE	Failure to Yield
58	PASSED ON RIGHT SHOULDER	Failure to Yield
59	PEDESTRIAN FAILED TO YIELD ROW TO VEHICLE	Failure to Yield
60	UNSAFE SPEED	Speeding
61	SPEEDING - (OVERLIMIT)	Speeding
62	TAKING MEDICATION	Impaired Driving
63	TURNED IMPROPERLY - CUT CORNER ON LEFT	Failure to Yield
64	TURNED IMPROPERLY - WIDE RIGHT	Failure to Yield
65	TURNED IMPROPERLY - WRONG LANE	Failure to Yield
66	TURNED WHEN UNSAFE	Failure to Yield
67	UNDER INFLUENCE - ALCOHOL	Impaired Driving
68	UNDER INFLUENCE - DRUG	Impaired Driving
69	WRONG SIDE - APPROACH OR INTERSECTION	Failure to Yield
70	WRONG SIDE - NOT PASSING	Failure to Yield
71	WRONG WAY - ONE WAY ROAD	Failure to Yield
72	CELL/MOBILE PHONE USE	Distracted Driving
73	ROAD RAGE	Other
74	OTHER (EXPLAIN IN NARRATIVE)	Other
75	CELL/MOBILE DEVICE USE - TALKING	Distracted Driving
76	CELL/MOBILE DEVICE USE - TEXTING	Distracted Driving
77	CELL/MOBILE DEVICE USE - OTHER	Distracted Driving
78	CELL/MOBILE DEVICE USE - UNKNOWN	Distracted Driving
79	FAILED TO SLOW FOR EMERGENCY LIGHTS	Failure to Yield
80	DROVE ON IMPROVED SHOULDER	Failure to Yield

Table 7: Mapping from original police contributing-factor codes to the 8-class taxonomy.



(a) Main labeling interface.



(b) Uncertainty marking view.

Figure 14: Annotation interface used for multi-label labeling and uncertainty marking.

# Prompts in the Wild: A Large Analyzed Collection of Transactional Prompts in Code

Victoria Basmov<sup>1,2</sup> Yoav Goldberg<sup>1,2</sup> Reut Tsarfaty<sup>1</sup>

<sup>1</sup>Bar-Ilan University <sup>2</sup>Allen Institute for Artificial Intelligence  
{vikasaeta, yoav.goldberg, reut.tsarfaty}@gmail.com

## Abstract

The behavior of contemporary generative Large Language Models (LLMs) is directly shaped by *prompts*, unstructured texts that describe the desired output and model behavior. In this paper we argue that prompts are linguistic objects that merit investigation in their own right. To this end, we collect 57.5K unique samples of prompts from GitHub. Specifically, we focus on transactional prompts: reproducible natural language instructions that are integrated into software. To enable the empirical, quantitative study of prompts, we introduce a structured ontology, capturing the properties of prompts as well as their formal and semantic components. Based on this ontology, we transform prompts from unstructured raw texts into richly structured linguistic objects. Analysis of these structured data reveals significant diversity of usage patterns across languages, domains, tasks, and modalities, in a typical Zipf-like distribution where some clearly prevail and others, more diverse, appear in the long tail. To validate the reliability of the ontology-based annotation of the prompts, we perform a comprehensive error analysis across all fields, providing a detailed assessment of annotation quality. We release the dataset together with a browsing and exploration interface.

## 1 Introduction

Prompts, the instructions humans give to large language models (LLMs), constitute the primary interface for guiding models' behavior. Despite growing practical interest in prompt engineering and prompting strategies, prompts are still treated largely as informal and intuitive artifacts rather than objects of systematic scientific inquiry. While models are analyzed in great depth, the usage of prompts, natural language utterances which directly shape models' behavior, remains largely ad-hoc (Villamizar et al., 2025). A systematic study of prompts may reveal crucial aspects of LLMs usage patterns: what

languages are used in prompts and how? what structural and semantic patterns do they follow? what tasks are they used to solve? what common practices emerge in prompt design? and a lot more. However, tapping into these questions and investigating them empirically, requires injecting into prompts structure that would allow for *quantitative, rigorous* analyses. Moreover, formalizing prompt structure is essential for the linguistic analysis of prompts (Jeoung et al., 2025; Leidinger et al., 2023); research of sensitivity to linguistic and structural prompt variation (Cuellar et al., 2026; Arabzadeh and Bagheri, 2025; Wahle et al., 2024); for tools and methods of structure-aware and linguistically informed automated prompt optimization (Santos et al., 2025; Hidalgo et al., 2025; Khattab et al., 2023; Saletta and Ferretti, 2024; Khattab et al., 2022; Murthy et al., 2025; Juneja et al., 2025); multilingual prompt engineering (Vatsal et al., 2025; Zhang et al., 2025; Kmainasi et al., 2024); and other areas of prompt research and downstream tasks.

We aim to establish prompts as first-order objects of scientific study. A phenomenon becomes a scientific object of study once practical relevance and sustained research are complemented by a shared *formal framework*, which prompts still lack. While prompts are already gaining interest, not only as tools for using LLMs but also as independent objects of study (Pister et al., 2024; Mao et al., 2025; Vir et al., 2025; Zheng et al., 2024a; Villamizar et al., 2025), prompt research lacks common terminology, structure, and large-scale empirical grounds. This work aims to fill this gap.

To facilitate the study of prompts, we collected a dataset of 57.5K unique prompts from public GitHub repositories. We specifically focus on *transactional prompts*,<sup>1</sup> prompts that are intended to

---

<sup>1</sup>The term "transactional prompts" is defined by M. Hashimoto (<https://mitchellh.com/writing/prompt-engineering-transactional-prompting>) to distinguish them from interactive prompts. They are also

perform reproducible, parameterized tasks, as part of a larger software-based workflow. Unlike casual (“interactive”) prompts, which are ad-hoc and one-off interactions with LLMs as part of a user-LLM conversation, transactional prompts run within pre-defined automated workflows and are refined to be robust and repeatable. Studying transactional prompts offers vital insights into real-world LLM usage in software applications.<sup>2</sup>

Unlike conventional programming, LLM instructions realized in prompts are specified in unstructured texts that convey complex, hierarchical, and multi-faceted messages, using natural language to encode a mixture of instructions and information. To explore the expressive power of this new “natural programming language” and the structural, compositional, and linguistic mechanisms through which this semantic range is expressed, it is helpful to introduce structure into otherwise unstructured prompt texts. To this end, we define an ontology underlying prompt structure.

The ontology provides a systematic framework for investigating the complex semantics encoded in prompts and the diversity of expressive means employed to express it (Section 3). It helps to uncover underlying semantic and structural patterns within widely diverse prompt data, providing a foundation for the empirical investigation of prompts as a “programming language” for LLMs across diverse applications, languages, and modalities.

Our analysis of the resulting structured prompts (Section 5) reveals a rich diversity of prompt usage across languages, modalities, tasks and domains, exhibiting a Zipf-like distribution — typical of linguistic phenomena (Piantadosi, 2014; Linders and Louwse, 2022) — with a prominent head and a much more diverse and nuanced long tail.

Importantly, our data collection and ontology are not intended to be final or exhaustive. Rather, they represent a starting point paving the way for fur-

---

sometimes referred to as “developer prompts”, but the latter term is often used in another sense: interactive prompts by software developers.

<sup>2</sup>Transactional prompts are predominantly single-turn. In contrast, multi-turn interactive user-LLM dialogues fall outside the scope of this work and are already addressed by existing large-scale datasets (Section 7). In addition, the emerging paradigm of agentic usage introduces prompts designed for iterative tool-use loops, which differ structurally and functionally from transactional prompts. Such prompts require dedicated investigation and potentially a specialized dataset, but they are beyond the scope of the present study. While our dataset may incidentally contain prompts that originally formed part of a chain or loop, they are represented as independent entries in our corpus.

ther investigation by linguists, prompt researchers, and engineers, who can contribute complementary perspectives to the research.

The prompt collection we deliver is accompanied by an online user interface for browsing, searching, and exploring prompts by their various characteristics and components, as defined by the underlying ontology. Alongside the collection and the ontology, it is intended as a practical resource to inspire further investigation into this topic.

In sum, this work treats prompts as first-class objects for empirical scientific and linguistic investigation, and makes four main contributions:

- (i) **a large-scale dataset** of 57.5K transactional prompts gathered from GitHub;
- (ii) **a structured prompt-ontology** that captures the primary prompt features and components;
- (iii) **empirical analysis** of the structured prompts, highlighting patterns in the way programmers use LLMs; and
- (iv) **a user interface** for browsing and searching prompts by their properties and components to support further research.

We aim for these resources to facilitate linguistic and practical investigations into how humans interact with LLMs and how prompts’ structure interacts with the semantic space they construct when using the prompting language as part of software engineering — pragmatic, syntactic, and lexical variations they use, how they structure prompts to elicit outputs with specific forms and content, the strategies they employ to overcome the inherent ambiguity and underspecification of natural language, how they choose the language of the prompt (English vs. other languages), are all prompts similar stylistically due to style convergence induced by LLMs, in what ways prompting language diverges from ordinary human language, and other potential perspectives.

## 2 Collecting Transactional Prompts

Following standard practices (Liu et al., 2025; Li et al., 2022; Mir et al., 2021; Alon et al., 2019; Raychev et al., 2016), we collect prompts from Github repositories<sup>3</sup>, by looking for files that either

---

<sup>3</sup>We use GitHub due to its availability and convenience. It is the largest repository of real world software projects, and its APIs allow flexible search and retrieval over the entire collection. We believe it is the current best source for locating prompts that are used as part of software projects. However, while large and diverse, it is also a biased source: for example, it does not include enterprise and closed-source projects, which may have a different distribution of prompts.

invoke the `chat.completion.create` API or the `PromptTemplate` constructor from the `LangChain` package<sup>4</sup>, and attempt to extract the contents of the `messages` (for `completion.create`) or `template` (for `PromptTemplate`) parameters from each call-site.

The immediate content is often a formatted string or a variable name, which we then aim to resolve to the actual prompt content via static analysis of the code, recursively tracking string values across variable assignments and function calls. The cases where this resolution succeeds are then filtered using a set of heuristics to retain only semantically-contentful prompts, and the filtered results are deduplicated. For each resulting prompt we further retain its metadata such as the repository name, file-path and URL, as well as the last commit date for the prompt. This process resulted in 57,640 unique prompts. Of these, 36,916 came from `chat.completion.create` and 20,724 from `PromptTemplate`. Details of prompt text extraction are available in the Appendix A. Details of filtering and deduplication are described in Appendix B.

### 3 The Prompts Ontology

In order to analyze prompts more deeply, we introduce an ontology outlining their main components and dimensions. Figure 1 shows a bird’s-eye view of our proposed ontology, with the specific fields explained shortly.

The ontology categories are grounded in three complementary sources: (1) Inherent prompt properties (e.g., prompts are texts written in a *language*; prompts by definition formulate a *task*. prompts serve to elicit certain *output*), (2) prior literature, e.g., prompting techniques (Jr et al., 2025; Schulhoff et al., 2025), context-grounded vs. parametric prompting (Zhou et al., 2024; Sun et al., 2026), and (3) recurring lexical and structural components we identified through manual inspection of sampled prompts, such as explicit language mentions, semantically distinct instruction blocks etc. The categories capture orthogonal dimensions of the data and are not intended to constitute a mutually exclusive or collectively exhaustive taxonomy. Rather, the ontology is deliberately non-restrictive and extensible; dataset users may adopt, adapt, extend, or

<sup>4</sup>We chose these two APIs due to their popularity and standardization: they are both widely adopted, and used in a consistent manner that make prompt extraction feasible. This comes at the expense of biasing the prompt collection to projects that use these APIs. This excludes, for instance, projects that write their own LLM access wrappers, or use other libraries.

replace it as appropriate for their purposes. The utility of the proposed ontology is directly demonstrated by its application in the data annotation and the analysis we present (Section 5).

**Languages:** detected languages used in prompt texts and any explicit language mentions.

**Task and domain:** we track the tasks for which the prompt is intended, and its application domain. These have both coarse-grained and fine-grained categories.

**Input characteristics:** At the outset, prompts specify one or more of (1) overall high-level instructions (“answer the question provided by the user”); (2) a question/task to be solved (“how many apples did John eat?”); (3) supporting context for 2. Each of these can be either hard-coded in the prompt or be a variable provided as input in each invocation. We identify the cases where each of these information types are read as input. For each identified input slot, we also retain information about its language, structure and modality, if available.

**Output characteristics:** Each prompt’s expected output is annotated for modality and for the requested structure, language and answer paradigm.

**Prompt structure:** Each prompt is represented as a list of role-messages (“System”, “User”, “Assistant”), and their associated texts.<sup>5</sup> For each message we list, beyond its text and role, also its detected language and languages explicitly mentioned within it. We also further break the message text into a sequence of individual instructions. Each instruction is associated with one of 42 semantic kinds (e.g. “role specification”, “audience specification”, “input content description”. See Appendix E for the complete list). We explicitly mark *negative* instructions, and distinguish between *central* and *auxiliary* instructions.

**Prompting techniques:** For each prompt we extract a list of prompting techniques with records specifying which techniques are used in the prompt, with supporting evidence spans. The prompting techniques come from a pre-specified list of 12 techniques (e.g. “use of sections”, “structured outputs”, see Figure 1).

<sup>5</sup>The `LangChain PromptTemplate` prompts are strings and not message sequences. These are represented as a single message with role “Undefined”.

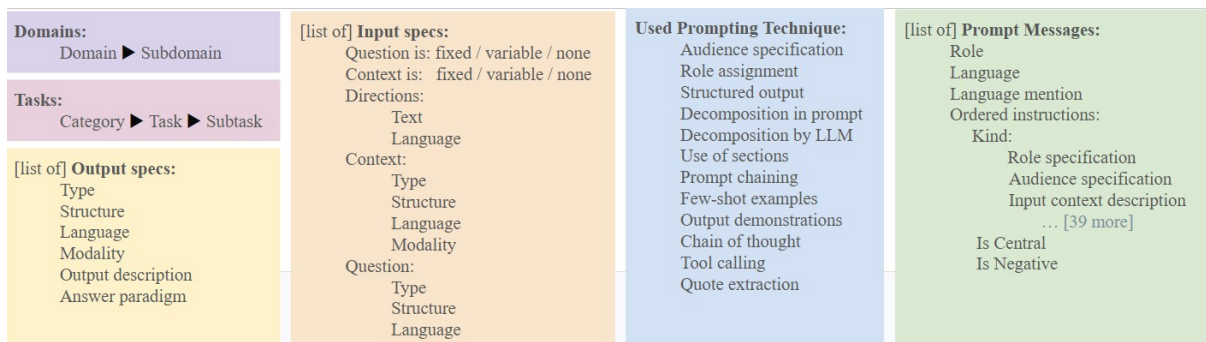


Figure 1: The **Prompt Ontology** Underlying the Empirical Analysis and the Structured Collection

**Meta-data:** Additionally, each object in the data includes an ID, a GitHub URL of the source file, timestamp of the last update, full prompt text (a concatenation of all individual message texts in the prompt) and its translation into English if not in English already.

For some fields (input and output structure, modality and variability, prompting techniques, instruction kinds, role), we defined the possible value inventories, although in some cases (e.g., structure fields) we instructed the model to enrich predefined values with additional detail (e.g., “Dictionary of items (‘From’: string, ‘To’: string)” rather than just “Dictionary”). For the remaining fields, values were generated by the LLM during annotation. For fields with especially large inventories (e.g., class, domain, input and output type), the values were then grouped into classes using the algorithm described in Appendix G.

#### 4 Annotation, Quality Control and Error Analysis

The ontology annotation is performed using an LLM-based process. The prompts used for annotating the prompt collection, and the detailed data-model, are listed in Appendix J.

We iteratively tested, manually evaluated and refined annotation prompts on sample data. Concretely, we jointly performed human-in-the-loop checks: we sampled  $\approx 100$  items per major meta-data category and manually inspected the automatic labels and their evidence spans. Any disagreements between the authors were resolved through peer discussions. These checks were used to refine the annotation prompts until labels became consistent, and the manual inspections showed consistently high agreement between the automatic labels and human judgments across categories.

To assess the resulting annotation quality, we

manually evaluate additional 100 randomly selected data points (50 from each source) and perform error analysis across all fields.<sup>6</sup> In the error analysis, predicted labels and evidence spans were evaluated against the annotation guidelines provided to the model in the respective prompts (see Appendix J). Detected errors were then grouped into recurring categories and quantified to characterize the main failure modes. The error analysis was conducted by a single expert, due to the substantial time required for this kind of fine-grained manual review.

The error analysis shows generally strong performance across most fields, with many categories exceeding 90% accuracy (e.g., Prompt Language 93.0%, Domain 90.6%, Context Language/Modality 97%, Central vs. Meta 96.4%, Prompting Techniques 98.5%). However, several fields remain more challenging, especially Output Type (60.4%) and Directions Text (69.4%), with moderate error rates in Answer Paradigm (80.2%), Instruction Kinds (89.8%), and Context Structure (81.9%). The observed errors fall into several recurring groups: (1) *hallucinations*, where the model assigns attributes not grounded in the prompt (e.g., inventing output types, domains, or languages); (2) *confusions between related categories*, such as conflating prompt language with explicit language mentions, labeling restrictions as negative instructions, or reporting output structure instead of context or question structure; (3) *omissions*, where relevant elements are not detected (e.g., missing tasks, ignored placeholders, undetected context or question units); and (4) *segmentation and granularity errors*, including failure to split complex instruction blocks or grouping multiple output types into a single “complex” type. Notably, many errors arise when indirect inference is required—e.g.,

<sup>6</sup>For detailed error analysis results see Tables 1–8 and Figure 5 in Appendix H.

coreference resolution, common-sense reasoning, or domain knowledge (inferring the output type, not mentioned directly, from few-shot examples, recognizing a Python function signature as code etc.). The results suggest that performance degrades primarily in cases involving implicit information and fine-grained annotation distinctions.

## 5 Analysis

### 5.1 Language, Modality and Domain

**Language.** Which languages are prompts written in? Naturally, English is predominant, but to what extent? And how diverse are the other languages? Our analysis shows that the dataset encompasses prompt messages in 62 languages. English is overwhelmingly dominant, accounting for 84.66% of identifiable cases.<sup>7</sup> The other languages that constitute above 1% are *Chinese, Korean, Spanish, Japanese and Portuguese*. The following seven highly represented languages are mostly (but not only) European: *French, Russian, German, Indonesian, Vietnamese, Polish, Italian and Dutch* (see Figure 6 Appendix I).<sup>8</sup> The long-tail languages occurring below 100 times in the data, with *Hindi* at the beginning and *Tagalog* at the very end, are shown in Figure 7 (Appendix I).

**Multilinguality.** Next, we examine the presence of multilinguism. Only 6.3% of prompts (3,649) are multilingual, and of these, over 99% include English. Despite this English dominance in mixed-language queries, 8.19% of the total dataset (4721 prompts) consists of entirely non-English text.

**Language Mentions.** Beyond the language used to write the prompt, we also annotate the prompts for explicit language mentions—instances where a language is referenced but not necessarily used (e.g., “Translate this to Afrikaans”). This often-overlooked dimension of multilinguality reveals even greater linguistic diversity. In terms of total mentions, we observe 15,331 references to natural human languages. In terms of language diversity, while English leads (9,510 mentions), the remaining references cover 151 languages and dialects, a far broader range than is found in the languages directly used in prompt texts. The listing and distribution of

<sup>7</sup>In 9.35% of the cases, the language (annotated per-span) could not be identified (e.g. when the span consisted of only a placeholder, like “{system\_txt}”).

<sup>8</sup>Interestingly, the highly represented languages include the top 10 prompt languages reported by Pister et al. (2024).

these non-English mentioned languages is detailed in Figure 8 (Appendix I).

**Modality.** While it is obvious that in LLMs text is the dominant input and output modality, what other combinations of input and output modalities do we see in the data? And can we find even greater diversity in the long tail? Where is the predominance of text more pronounced: in input or in output?

In the overwhelming majority of cases, the input context modality (77.82% of the cases) and the output modality (over 97% of the cases) is text.<sup>9</sup> The distributions of non-text or mixed modality for input and output are shown on the Figures 9 and 10 (Appendix I) with images accounting for a much greater share than audio and video.

The frequent input-output modality combinations all have textual output and differ only in input: the prevailing combination is text→text (77.31% of all the data); with the other categories—ungrounded or undefined (18.98%), image→text (1.98%), image+text→text (0.66%), audio→text (0.23%) — lag far behind (Fig 11, Appendix I).

**Domain.** In which domains are transactional prompts most frequently used? Which domains are leading and which are in the long tail? Prompts with a specific domain that can be identified<sup>10</sup> constitute 57.38% of the prompts. This resulted in 77 distinct domains with a long-tail, Zipf-like distribution. The top leading domains are, in order: *education & instruction, software development, business & commerce, healthcare & medical, technology, media & entertainment, finance & banking, creative writing & content-creation, human resources, arts & culture*. Together, they cover almost half (49.58%) of all the domains in the data, and appear in 63.63% of the prompts with specific domains.<sup>11</sup> Mid-frequency domains include, for example, *hospitality & food service, design & arts, personal services* and *philosophy*, while low-frequency domains include *urban development, historical studies, politics* and others. See Appendix D for the

<sup>9</sup>The remaining cases are those where no context was identified (10.75% of the prompts) or the input or output modality was unidentifiable (for example, the input modality can remain undefined when the input is a PDF document, but the prompt does not specify if it contains text, an image, or both). This covers 8.23% of the inputs and 2.77% of the outputs.

<sup>10</sup>Cases where a domain can not be identified are often short general prompts such as, “Please answer the user question using only the given context.”

<sup>11</sup>A single prompt can belong to multiple domains.

full list of domains.

## 5.2 Structure and Semantics

**Instruction Kinds.** A prompt (or a message therein) can be interpreted as a sequence of instructions given to the LLM. But what is the semantic or functional structure of prompts? That is, what are the semantic types of instructions and their order in transactional prompts? Our ontological structure treats each message as a sequence of instruction items, each labeled with its semantic function.

Overall, the dataset includes 39,4875 such instruction blocks, 6.85 per prompt on average. The top 10 most frequent types are:

- Input context placeholder (16%)
- Constraint or restriction (11%)
- Output content requirements (9.7%)
- Output format requirement (9.7%)
- Role specification (8.1%)
- Input context description (7.1%)
- Central task/question description (6.7%)
- Central task/question (5.8%)
- Input contextual data (4.3%)
- Conditional instruction (2.9%)

Together, they cover 81.33% of all blocks in the data. Full statistics, as well as information about frequent ordering of units, are available in Appendix I, Figure 18, and Tables 10 & 11.

**Core vs. Supporting Instructions.** Generally, 18.2% of instructions represent the Central Task (the “core intent”), while the remaining 81.8% function as meta (supporting) instructions providing guidance on style, constraints, or formatting. Most of the core instructions focus on the task/question detailed descriptions (36%) or define the tasks/questions themselves (31.9%).

In the example below, the core instruction defines *what* the task is while the supporting ones add details by specifying *how* it should be performed: "Please create a learning plan in {language}." (*core*) "The plan should outline daily activities." (*supporting*). "Make sure to include detailed information about the specific programming languages and tools (like APIs) that will be used." (*supporting*) "Do not include learning of languages that I have already used." (*supporting*).

Only 4% of prompts consist exclusively of a central task with no metadata. These are typically short, non-grounded queries (e.g., “Explain how to write a window in Python”, “Name ten mammals”) or where the context is not provided in the prompt

text (e.g., “Identify products using the given images and generate key features for each product.”).

**Negative Instructions.** Negative instructions (telling the model what not to do) serve as a window into user mitigation of LLMs tendencies, such as verbosity, hallucination and different kinds of bias.

Roughly 31% of all prompts contain at least one negative instruction. Overall, negative instructions represent only 7% of the total instruction count in the dataset. The vast majority (89.5%) of the negative instructions are categorized as constraints or restrictions. These act as guardrails against undesirable behaviors, such as adding extra text beyond the requested output, or exceeding specific lengths. A smaller portion addresses output format (5.4%), content requirements (1.4%), and error handling (1.1%). Examples of negative instructions include (more in Appendix F):

- *don't make the answers too long* (constraint/restriction)
  - *If you encounter an exception, an effectless command, or find yourself in a loop, avoid repeating the same command and try something else to achieve the goal.* (error handling instruction)
  - *Don't translate the text to English. Keep it in Indonesian.* (linguistic constraint/specification)
- Negative instructions very rarely form the core intent of the prompt (0.5% of cases). In these instances, the primary task is defined by what the model must avoid, e.g., “Do not answer the questions, simply provide a correct compute graph...” or “Do not respond to text, merely translate it.”

**Constraints.** We define constraints as instructions involving restrictions, style/format requirements, design specifications, or negative directives. As research shifts toward evaluating how well LLMs adhere to nuanced requirements (Zhou et al., 2023; Lior et al., 2025), our dataset emerges as a source of naturally occurring constraints. Like negative instructions, the isolation of prompt constraints also offers opportunities for linguistic investigations on the form and structure in which they are expressed.

Our analysis reveals a landscape of high complexity. Constraints represent 33.3% of all instruction blocks in the dataset. On average, a single prompt contains 2.28 constraints, but the “long tail” is significant — some transactional prompts contain up to 155 distinct rules. Furthermore, while 27.4% of the prompts are constraint-free, nearly 14% layer five or more constraints in a single query, signaling a demand for high-precision model control.

Unlike synthetic benchmarks, our data captures the “messy” and layered constraints actually deployed in real-world scenarios, making it a unique resource for investigating the limits of LLM steerability in real-world transactional environments.

**Messages.** The popular `chat.completions.create` API expects prompts as a sequence of role messages (“System”, “User”, “Assistant”). How do prompt writers use this interface? How many messages do they use, and what roles are assigned? The majority of prompts (64%) have two messages. Of these 96% are “system-user”. 27% of the prompts have a single message, with “user” (70%) being far more frequent than “system” (28%). For three-message prompts (4%) the most popular sequences are “system-user-user” (26.41%), “system-user-assistant” (22.01%), “system-system-user” (18.45%), “system-assistant-user” (13.07%). For prompts with more than three messages, the majority (52%) includes alternations of the “user” and “assistant” messages, alternatively with a system prompt at the beginning. (See details in Figure 16 Appendix I and Figure 17 Appendix I). It thus appears that the “system-user” duo has become the standard unit in transactional prompts. This suggests that developers largely view the System role as a static configuration layer and the User role as the dynamic input, rather than utilizing the API for complex, multi-turn role-play within a single template.

### 5.3 What Are Prompts Being Used For?

**Tasks.** What tasks are typically performed using LLMs? Are LLMs used more for standard NLP tasks or for other, non-NLP tasks? From these tasks (NLP vs. non-NLP), to what extent either is more pronounced? Are the tasks used across domains or are some of them limited to a specific domain? The distribution of NLP tasks in the data covers both NLP tasks (involving language understanding, generation and analysis) and tasks outside classical NLP (like data processing, analysis, multimodal, and structured-data tasks).

It is clear from the data that NLP tasks prevail: the top 4 (question answering, general text generation, information extraction, and summarization) cover over 48% of the data (see Figure 13, Appendix I). Non-NLP tasks occur mostly at moderate to low frequencies. In the long tail (tasks with under 30 cases) we see non-NLP tasks, such as education design, game strategy, state tracking, data privacy,

policy generation, and system integration. The top 10 in long-tail and mid-frequency tasks are shown in Table 9 (Appendix I). We further see that each task appears in multiple domains, from 4 (game strategy) to 77 (question answering), and no task is purely domain-specific.

**Inputs.** It is common for prompts to have a *question* or *main task* that needs to be answered, either based on a *context* that is also provided, or based on parametric knowledge. Either of these components can be hard-coded into the prompt, or be a varying input to the prompt. In our data, 97.9% of the prompts included a question or a main task, and 89.2% included a supporting context. From these, the question/task is expected as input 70% of the time, and was hard-coded in the prompt for the remaining 30%. In contrast, a context is provided as input 94% of the time and is only hard coded 6% of the time. This suggests an (expected) tendency to perform the same task over varying contexts rather than performing varying tasks over a fixed context.

**Grounding.** LLMs may be expected to respond either based on their internal parametric knowledge, or based on grounding context provided to them. Which of these options is more prevalent in real-world transactional usage? And, in the case of grounded prompts, what kinds of contexts are used for grounding?

Our analysis suggests that 89.25% of all cases are grounded, i.e., performed based on a certain context rather than merely based on the LLM’s parametric knowledge.<sup>12</sup> In the vast majority of cases the type of the grounding context is text (that may include a variety of subtypes such as document, paragraph, sentence, abstract, tweet, proverb, caption, description, instruction, etc.). Other frequent types of context include a question, dialogue/conversation, code, table, image, numeric context, json. The distribution of top 10 input types across the top 10 tasks is shown in Figure 14 (Appendix I).

### 5.4 Prompting Techniques

Which prompting techniques are adopted by prompt writers, and to what extent?

*Role assignment* is by far the most frequently used technique, accounting for over 45% of all instances

<sup>12</sup>The cases where no context was found are not necessarily non-grounded. This can be due to other reasons. Sometimes the context is added to the prompt dynamically, beyond the `chat.completions.create` API call or the `PromptTemplate` static instantiation, and in these cases our system was unable to capture it.

(Figure 19, Appendix I). Explicitly defining the role of the AI was one of the first well-known techniques, and was recommended by model developers. The data shows that this remains a widely adopted practice in prompt engineering, despite reports of its limited effectiveness in more recent models (Zheng et al., 2024b; Kim et al., 2025).

The next most commonly-used technique is *structured output* (12.9%), that is, requesting output in machine parseable formats (json, XML etc.). This technique is closely related to the *output demonstrations* technique (3.34%) including examples of how the output should be organized. Together, these constitute 16.24% of the cases. Also related is the *sections* technique (7.83%) dividing the prompt into clearly defined, marked sections. These demonstrate the significance of clearly-structured input and output specifications.

The next in frequency is *decomposition via prompt* (10.34%) meaning that the prompt specifies how to break the task down into sub-tasks or manageable steps. This technique is complemented by the much less frequent *decomposition via LLM* (0.56%) where the same decomposition is expected to be done by the LLM. Together, these techniques form 10.9% of the annotated techniques, suggesting the importance of solving complex tasks by breaking them down into more manageable units.

## 6 User Interface

To empower researchers to explore the prompt collection beyond our set of analyses, we provide a web-based UI designed for exploratory discovery and deep-dive analyses into subsets of the data. Users can filter prompts by ontology fields; search the dataset with semantic similarity; see and aggregate counts; inspect individual prompts and their annotated sections; and download prompts and ontological data for their filtered subsets. Additional information on the UI can be found in Appendix C.

## 7 Related Work

As interest in prompts as distinctly designed artifacts has grown, several datasets have emerged, encompassing both *user prompts* and *transactional prompts*.

**User prompts** are prompts that are intended to be used directly by users, and have significantly different characteristics than the transactional prompts we study herein. LMSYS-Chat-1M (Zheng et al., 2024a) is a dataset of 1M user-LLM conversations collected across 154 languages via the Vicuna demo

and Chatbot Arena. WildChat (Zhao et al., 2024) is a multilingual, dataset of 1 million timestamped user-LLM conversations (over 2.5 million turns) collected via a ChatGPT/GPT-4 chatbot with explicit user consent, annotated with demographic metadata (state, country, and hashed IPs) to enable behavioral analysis. PROMPTEVALS (Vir et al., 2025) is a dataset of 2,087 LLM prompt templates from the LangChain Prompt Hub,<sup>13</sup> a repository of user-contributed prompts, containing a mix of user-prompts and transactional prompts. PROMPTEVALS is intended for training and evaluating “assertion guardrails”, and has prompts spanning multiple domains including IT, finance, and healthcare. DevGPT (Xiao et al., 2023) is a dataset of 29,778 ChatGPT prompts and responses from software developers, collected from shared ChatGPT conversations on GitHub and Hacker News for analysis of developers’ interactions with ChatGPT and their implications for AI-assisted programming. These datasets are different than ours in that they do not address transactional prompts in software.

Other researchers have studied **Non-LLM prompts**, for instance DiffusionDB (Wang et al., 2023) and VidProM (Wang and Yang, 2024) that compile large-scale prompts for text-to-image and text-to-video generation, respectively. As proposed herein, such collection too would merit from structured representation and empirical analysis, compatible to ours, which is reserved for future research.

Finally, for **Transactional Prompts**, Pister et al. (2024) introduced PromptSet, a dataset of developer prompts with similar size and collection method to our own. However, they invest less effort in extraction and cleanup compared to us. As a result, as reported by Tafreshipour et al. (2025) and verified by us, the data contains many incomplete prompts or prompt fragments that are hard or impossible to analyze. They also do not provide analysis or structuring of the prompts beyond this raw string data. In spite of these limitations, researchers use PromptSet for exploration of different aspects of transactional prompts (Villamizar et al., 2025; Tafreshipour et al., 2025; Mao et al., 2025), as well as for development of prompt optimization tools (Rzig et al., 2025). Notably, Mao et al. (2025) construct their own small dataset, derived from PromptSet, to analyze real-world prompts that combine static content with dynamic placeholders such as “input”. After filtering, cleaning and deduplication, they

<sup>13</sup><https://smith.langchain.com/hub>

extract 2,163 such prompts, in which they identified key components and categorized them into one of six semantic categories.

This highlights the interest in transactional prompts research and a clear community demand for larger, higher-quality resources for such research like the corpus presented in this work.

## 8 Conclusions

We present a large, real-world collection of transactional prompts; an ontology that captures both the structural components of prompts and their descriptive characteristics; and a web interface for their systematic exploration. These resources enable a range of applications, including linguistic and structural analysis of prompt texts, uncovering common conventions and “unspoken norms” of prompt composition, comparing recommended versus actual practices, and supporting multiple downstream applications, such as instruction-following and prompt-sensitivity research, more realistic benchmarking, structure-aware and linguistically-informed automated prompt optimization, multilingual prompt engineering and others. We present a preliminary empirical analysis that exemplifies the utility of the provided framework and illustrates some of the kinds of insights this data makes possible, and hope it inspires further interest in the systematic and methodological study of prompts as scientific and linguistic objects in the community.

## Limitations

Our search for prompts in GitHub files was limited to certain patterns (the `chat.completion.create` API call and the `LangChain PromptTemplate` class), whereas many other patterns are possible. Furthermore, we only considered Python files. Thus, we do not present an unbiased sample of transactional prompts, and the observed trends in our analysis may be biased towards a subset of the prompt space defined by users who chose to use these APIs. Additionally, while prompts are continuously created and updated, our dataset represents a snapshot at the time of collection. Therefore, rather than being exhaustive, our work paves the way to future efforts that could expand the collection, potentially making it dynamic by regularly incorporating new entries over time, and include broader search patterns and additional programming languages.

The prompt structuring annotations were performed by an LLM, and, though annotation prompts

were iteratively tested, manually evaluated and refined on sample data, each annotation result could not be manually verified individually. As a result, the data may still contain some errors or noise despite the efforts to ensure annotation quality.

Finally, the analysis presented in this work only scratches the surface, and many more quantitative and qualitative investigations are possible, including detailed linguistic analysis, diachronic comparisons (changes over time), cross-lingual or domain and task-specific exploration. We encourage the community to contribute to expanding the data and improving annotation accuracy and to continue and deepen the research in the field of transactional prompts.

## Acknowledgments

Work on this project was supported by a VATAT grant from the Planning and Budgeting Committee of the Council for Higher Education in Israel, Kamin grant by the Israel Innovation Authority (IIA) and ISF grant number 670/23.

## References

- Uri Alon, Shaked Brody, Omer Levy, and Eran Yahav. 2019. [code2seq: Generating sequences from structured representations of code](#). In [International Conference on Learning Representations](#).
- Negar Arabzadeh and Ebrahim Bagheri. 2025. [VAP3: Variation-aware prompt performance prediction](#). [Proceedings of the 48th International ACM SIGIR Conference on Research and Development in Information Retrieval](#).
- Jaime E. Cuellar, Óscar Moreno-Martínez, Paula Sofía Torres Rodríguez, Jaime Andrés Pavlich-Mariscal, Andrés Felipe Micán Castiblanco, and Juan Guillermo Torres Hurtado. 2026. [Trusting ChatGPT? When a subtle variation in the prompt can significantly alter the results](#). [Journal of Artificial Intelligence and Technology](#).
- Nicolás Hidalgo, Pablo Alzaga Sáez, Nicolas Meneeses, Víctor Reyes, and Erika Rosas. 2025. [Prompt’s evolution for language model-driven data generation](#). [Applied Sciences](#).
- Sullam Jeoung, Yueyan Chen, Yi Zhang, Shuai Wang, Haibo Ding, and Lin Lee Cheong. 2025. [PromptPrism: A linguistically-inspired taxonomy for prompts](#). [ArXiv, abs/2505.12592](#).
- E. G. Santana Jr, Gabriel Benjamin, Melissa Araujo, Harrison Santos, David Freitas, Eduardo Almeida, Paulo Anselmo da M. S. Neto, Jiawei Li, Jina Chun, and Iftekhar Ahmed. 2025. [Which prompting technique should i use? An empirical investigation of](#)

- [prompting techniques for software engineering tasks](#). Preprint, arXiv:2506.05614.
- Gurusha Juneja, Gautam Jajoo, Nagarajan Natarajan, Hua Li, Jian Jiao, and Amit Sharma. 2025. [Task facet learning: A structured approach to prompt optimization](#). Preprint, arXiv:2406.10504.
- Omar Khattab, Keshav Santhanam, Xiang Lisa Li, David Hall, Percy Liang, Christopher Potts, and Matei Zaharia. 2022. [Demonstrate-search-predict: Composing retrieval and language models for knowledge-intensive NLP](#). arXiv preprint arXiv:2212.14024.
- Omar Khattab, Arnav Singhvi, Paridhi Maheshwari, Zhiyuan Zhang, Keshav Santhanam, Sri Vardhamanan, Saiful Haq, Ashutosh Sharma, Thomas T. Joshi, Hanna Moazam, Heather Miller, Matei Zaharia, and Christopher Potts. 2023. [DSPy: Compiling declarative language model calls into self-improving pipelines](#). Preprint, arXiv:2310.03714.
- Junseok Kim, Nakyeong Yang, and Kyomin Jung. 2025. [Persona is a double-edged sword: Rethinking the impact of role-play prompts in zero-shot reasoning tasks](#). In [Proceedings of the 14th International Joint Conference on Natural Language Processing and the 4th Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics](#), pages 848–862, Mumbai, India. The Asian Federation of Natural Language Processing and The Association for Computational Linguistics.
- Mohamed Bayan Kmainasi, Rakif Khan, Ali Ez-zat Shahroor, Boushra Bendou, Maram Hasanain, and Firoj Alam. 2024. [Native vs non-native language prompting: A comparative analysis](#). ArXiv, abs/2409.07054.
- Alina Leidinger, Robert van Rooij, and Ekaterina Shutova. 2023. [The language of prompting: What linguistic properties make a prompt successful?](#) ArXiv, abs/2311.01967.
- Zhiyu Li, Shuai Lu, Daya Guo, Nan Duan, Shailesh Jannu, Grant Jenks, Deep Majumder, Jared Green, Alexey Svyatkovskiy, Shengyu Fu, and Neel Sundaresan. 2022. [Automating code review activities by large-scale pre-training](#). [Proceedings of the 30th ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering](#).
- Guido Linders and Max Louwerse. 2022. [Zipf’s law revisited: Spoken dialog, linguistic units, parameters, and the principle of least effort](#). [Psychonomic Bulletin Review](#), 30.
- Gili Lior, Asaf Yehudai, Ariel Gera, and Liat Ein-Dor. 2025. [WildIFEval: Instruction following in the wild](#). Preprint, arXiv:2503.06573.
- Jingjing Liu, Zeming Liu, Zihao Cheng, Mengliang He, Xiaoming Shi, Yuhang Guo, Xiangrong Zhu, Yuanfang Guo, Yunhong Wang, and Haifeng Wang. 2025. [RepoDebug: Repository-level multi-task and multi-language debugging evaluation of large language models](#). Preprint, arXiv:2509.04078.
- Yuetian Mao, Junjie He, and Chunyang Chen. 2025. [From prompts to templates: A systematic prompt template analysis for real-world LLMapps](#). Preprint, arXiv:2504.02052.
- A. M. Mir, E. Latoskinas, and G. Gousios. 2021. [Many-Types4Py: A benchmark python dataset for machine learning-based type inference](#). In [IEEE/ACM 18th International Conference on Mining Software Repositories \(MSR\)](#), pages 585–589. IEEE Computer Society.
- Rithesh Murthy, Ming Zhu, Liangwei Yang, Jieli Qiu, Juntao Tan, Shelby Heinecke, Caiming Xiong, Silvio Savarese, and Huan Wang. 2025. [Promptomatix: An automatic prompt optimization framework for large language models](#). Preprint, arXiv:2507.14241.
- Steven T. Piantadosi. 2014. [Zipf’s word frequency law in natural language: A critical review and future directions](#). [Psychonomic Bulletin & Review](#), 21(5):1112–1130.
- Kaiser Pister, Dhruva Jyoti Paul, Ishan Joshi, and Patrick Brophy. 2024. [PromptSet: A programmer’s prompting dataset](#). In [Proceedings of the 1st International Workshop on Large Language Models for Code, LLM4Code ’24](#), page 62–69. ACM.
- Veselin Raychev, Pavol Bielik, and Martin Vechev. 2016. [Probabilistic model for code with decision trees](#). pages 731–747.
- Dhia Elhaq Rzig, Dhruva Jyoti Paul, Kaiser Pister, Jordan Henkel, and Foyzul Hassan. 2025. [An empirically-grounded tool for automatic prompt linting and repair: A case study on bias, vulnerability, and optimization in developer prompts](#). ArXiv, abs/2501.12521.
- Martina Saletta and Claudio Ferretti. 2024. [Exploring the prompt space of large language models through evolutionary sampling](#). [Proceedings of the Genetic and Evolutionary Computation Conference](#).
- Gabriel Machado Santos, Rita Maria Silva Julia, and Marcelo Zanchetta do Nascimento. 2025. [Diverse prompts: Illuminating the prompt space of large language models with MAP-Elites](#). [2025 IEEE Congress on Evolutionary Computation \(CEC\)](#), pages 1–8.
- Sander Schulhoff, Michael Ilie, Nishant Balepur, Konstantine Kahadze, Amanda Liu, Chenglei Si, Yin-heng Li, Aayush Gupta, HyoJung Han, Sevien Schulhoff, Pranav Sandeep Dulepet, Saurav Vidyadhara, Dayeon Ki, Sweta Agrawal, Chau Pham, Gerson Kroiz, Feileen Li, Hudson Tao, Ashay Srivastava, and 12 others. 2025. [The prompt report: A systematic survey of prompt engineering techniques](#). Preprint, arXiv:2406.06608.

- Kaiser Sun, Fan Bai, and Mark Dredze. 2026. [Task matters: Knowledge requirements shape LLM responses to context-memory conflict](#). [Preprint](#), arXiv:2506.06485.
- Mahan Tafreshipour, Aaron Imani, Eric Huang, Eduardo Almeida, Thomas Zimmermann, and Iftekhar Ahmed. 2025. [Prompting in the wild: An empirical study of prompt evolution in software repositories](#). [Preprint](#), arXiv:2412.17298.
- Shubham Vatsal, Harsh Dubey, and Aditi Singh. 2025. [Multilingual prompt engineering in large language models: A survey across NLP tasks](#). [ArXiv](#), abs/2505.11665.
- Hugo Villamizar, Jannik Fischbach, Alexander Korn, Andreas Vogelsang, and Daniel Mendez. 2025. [Prompts as software engineering artifacts: A research agenda and preliminary findings](#). [Preprint](#), arXiv:2509.17548.
- Reya Vir, Shreya Shankar, Harrison Chase, Will Fu-Hinthorn, and Aditya Parameswaran. 2025. [PROMPTEVALS: A dataset of assertions and guardrails for custom production large language model pipelines](#). [Preprint](#), arXiv:2504.14738.
- Jan Philip Wahle, Terry Ruas, Yang Xu, and Bela Gipp. 2024. [Paraphrase types elicit prompt engineering capabilities](#). [ArXiv](#), abs/2406.19898.
- Wenhao Wang and Yi Yang. 2024. [VidProM: A million-scale real prompt-gallery dataset for text-to-video diffusion models](#). [Preprint](#), arXiv:2403.06098.
- Zijie J. Wang, Evan Montoya, David Munechika, Haoyang Yang, Benjamin Hoover, and Duen Horng Chau. 2023. [DiffusionDB: A large-scale prompt gallery dataset for text-to-image generative models](#). [Preprint](#), arXiv:2210.14896.
- Tao Xiao, Christoph Treude, Hideaki Hata, and Kenichi Matsumoto. 2023. [DevGPT: Studying developer-ChatGPT conversations](#). 2024 [IEEE/ACM 21st International Conference on Mining Software Repositories \(MSR\)](#), pages 227–230.
- Lechen Zhang, Yusheng Zhou, Tolga Ergen, Lajanugen Logeswaran, Moontae Lee, and David Jurgens. 2025. [Cross-lingual prompt steerability: Towards accurate and robust LLM behavior across languages](#). [ArXiv](#), abs/2512.02841.
- Wenting Zhao, Xiang Ren, Jack Hessel, Claire Cardie, Yejin Choi, and Yuntian Deng. 2024. [WildChat: 1M ChatGPT interaction logs in the wild](#). [Preprint](#), arXiv:2405.01470.
- Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Tianle Li, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zhuohan Li, Zi Lin, Eric P. Xing, Joseph E. Gonzalez, Ion Stoica, and Hao Zhang. 2024a. [LMSYS-Chat-1M: A large-scale real-world LLM conversation dataset](#). [Preprint](#), arXiv:2309.11998.
- Mingqian Zheng, Jiabin Pei, Lajanugen Logeswaran, Moontae Lee, and David Jurgens. 2024b. [When "a helpful assistant" is not really helpful: Personas in system prompts do not improve performances of large language models](#). [Preprint](#), arXiv:2311.10054.
- Jeffrey Zhou, Tianjian Lu, Swaroop Mishra, Siddhartha Brahma, Sujoy Basu, Yi Luan, Denny Zhou, and Le Hou. 2023. [Instruction-following evaluation for large language models](#). [Preprint](#), arXiv:2311.07911.
- Sizhe Zhou, Sha Li, Yu Meng, Yizhu Jiao, Heng Ji, and Jiawei Han. 2024. [Establishing knowledge preference in language models](#). [ArXiv](#), abs/2407.13048.

## A Details of GitHub Prompt Collection

To identify Python files containing prompts, we used the GitHub Code Search API to systematically query repositories for code involving prompt-related functions. Specifically, we searched for files that invoked `chat.completions.create`, a common method used in prompt construction for language models, and the `langchain PromptTemplate` class, a class used to generate prompts from a string template and variables. For each search result we collected metadata such as repository name, file path, and URL. This way we end up with 95806 objects from 95434 filepaths from 51393 repositories.

Building on our initial URL collection, we then implement an extraction pipeline that pulls actual prompt text out of each discovered file. We iterate through each file record, and retrieve file contents via GitHub REST API, decoding the Base64 encoded result into plain Python source code. We then parse that source with Python's `ast` module to locate all occurrences of our target API call, `chat.completions.create`, or the `langchain PromptTemplate` class. Then we employ a multistep process aiming to extract the full contents of the "messages" or "template" parameter (for `chat.completions.create` or `PromptTemplate` respectively) - even when they're built up across several statements. Specifically, using the `ast` module, we track variable assignments and resolve all arguments and keyword values used within the API call (whenever possible). If the API call is inside a function, we find where that function is called and replace its parameters with the actual values passed into the function at each call site - using both the current file and related imports.

Next, we check for remaining unresolved variables. If the entire messages field or a specific content field inside a messages list is a variable placeholder, the actual values of these variables are looked up in the current file and other related files in the repository. At every step, if a variable is reassigned to different values before different calls, our extraction logic will capture each distinct value, yielding multiple versions of the prompt.

Figure 2 illustrates some stages of this process.

Finally, for each prompt, we look up the date of most recent commit that changed any of the lines contributing to it, in order to estimate when the prompt was last modified. To ensure correctness of the extraction pipeline, we manually evaluated a subset of 1,000 prompts before scaling to the full

dataset. This extraction process leaves us initially with 145553 objects.

Next we perform filtering and deduplication (see Appendix B for details).

## B Filtering And Deduplication Details

We filter out objects where the extracted texts are empty, contain invalid values (e.g., 'error', 'n/a', 'nan'), or consist only of unresolved variables or placeholders, identifiable via string matching or regular expressions. Next, we apply a series of additional heuristics to filter out prompts that lack readable or meaningful content. Specifically, we remove prompts that consist solely of punctuation or whose language cannot be reliably detected by the `langdetect` library. For prompts identified as English, we use `spaCy` to parse the text and check for the presence of verbs or auxiliaries—signals of syntactic structure and potential informativeness. Prompts with such features are retained. Prompts in clearly detected non-English languages are also kept. These heuristics help exclude most empty, malformed, or placeholder-based prompts while preserving those that exhibit valid language or meaningful structure.

The deduplication procedure is as follows. Exact repeats, defined as objects with the same file path, extracted prompt text, and timestamp, are removed after the first instance. Prompt texts that occur more than once but in different files or at different times are retained, but marked as duplicates. This yields a dataset of 85,209 objects. In this version of the dataset, we identify 8,169 groups of duplicate prompts. Group sizes range from 2 to 580 prompts. Most groups (63.55%) contain only 2 duplicate prompt instances, followed by 15.79% with 3 instances, 6.24% with 4 instances, 7.28% with 5–7 instances, and the remaining 7.14% with 8 or more instances.<sup>14</sup>

Additionally, we create a fully deduplicated version of the dataset in which cross-file repeats are removed after the first instance. The analysis and statistics reported in this article are based on the

---

<sup>14</sup>The most frequently reproduced prompt, appearing 580 times, is:

*Answer the question as detailed as possible from the provided context, make sure to provide all the details. If the answer is not in the provided context, just say, 'answer is not available in the context' and do not provide a wrong answer.*

*Context: {context}*

*Question: {question}*

*Answer:*

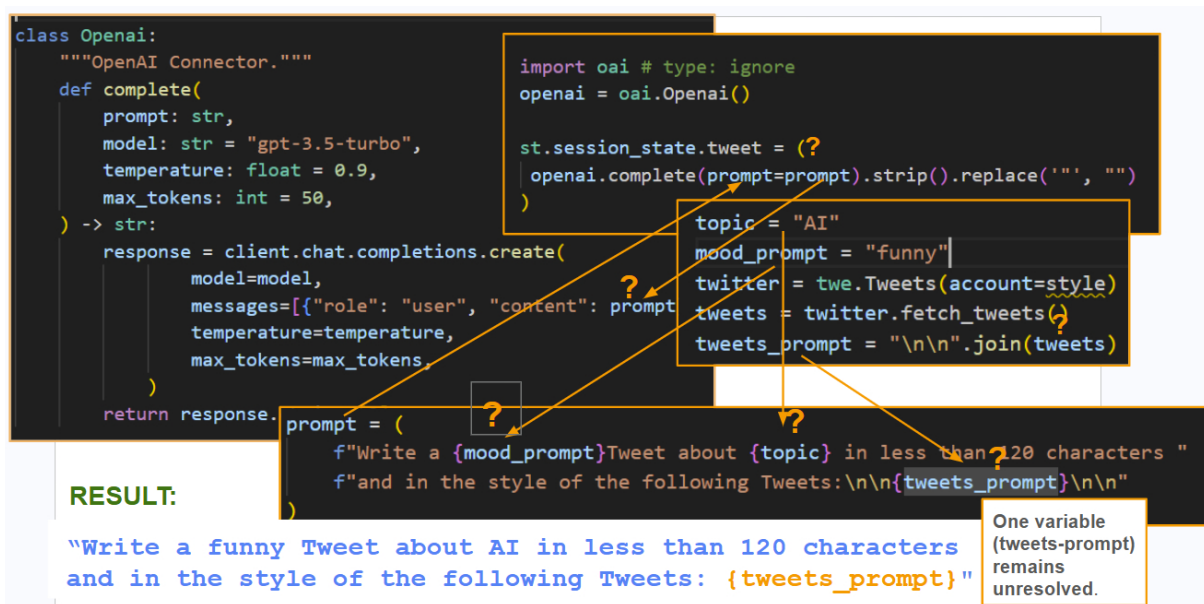


Figure 2: Prompt Extraction Flow. This figure illustrates prompt text extraction by tracing variables across the repository. Static variables (such as `topic` and `mood_prompt`) are successfully resolved, while dynamic variables requiring runtime execution (such as `tweets_prompt`) remain unresolved in the final extracted text.

fully deduplicated dataset version.<sup>15</sup>

## C User Interface: Layout and Functionality

The layout and functionality of each UI component are as follows. At the top are a semantic free-text search field, a filter box showing all active filters, *Show prompts* and *Download prompts* buttons. Below them, the page displays a set of boxes for different ontology fields (task, domain, language, modality, prompting-techniques, etc.). Each box lists all available values for the field with the corresponding prompt counts, which update dynamically as filters are applied. It shows coarse categories by default. Clicking an eye icon next to each coarse category reveals its fine-grained subcategories.

Users can select multiple values in each box and switch between *match-all* and *match-any* mode. Clicking the checkbox next to a value selects or deselects it. The "Apply filters" button applies the selected filters. The filter box offers per-filter removal and a *Clear all filters* button. Counts across all boxes update dynamically as filters change.

The *Show prompts* button opens a paginated drawer containing a stack of prompt cards. Each card displays the prompt text and includes a *Show spans* toggle that marks structural components - di-

rections, context, question/task, output description and different semantic kinds of instruction blocks - using colors. A color legend on each card explains the span colors. Hovering a legend entry highlights the corresponding spans in the prompt for convenience.

Free-text search uses embedding-based semantic similarity: the system embeds each query with the embedding `gemma-300m-ONNX` model, which is also used to precompute embeddings for all the prompts, and retrieves relevant prompts by cosine similarity.

The header displays the number of prompts matching the current filters. Clicking the *Download prompts* button exports the full dataset entries for the selected prompts.

Figures 3 and 4 illustrate some features of the user interface.

## D Domains

Below all the domains in our data are listed along with their number and percentage.

1. education & instruction - 4182 (8.42%)
2. software development - 3863 (7.78%)
3. business & commerce - 2790 (5.62%)
4. healthcare & medical - 2485 (5.00%)
5. technology - 2468 (4.97%)
6. media & entertainment - 2040 (4.11%)

<sup>15</sup>Indeed, patterns and reasons for prompt reuse may be of interest for future analysis; however, in this work, we focus on the structure and diversity of prompt design .

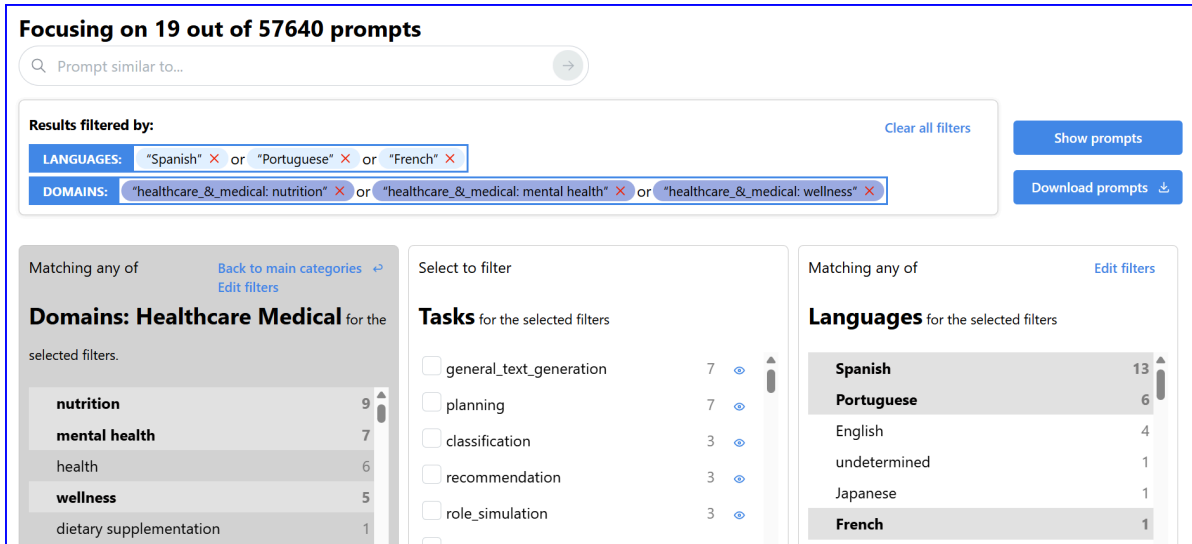


Figure 3: User Interface. The top section features a free-text search field, a filter box displaying currently active filters, and buttons for prompt display and download. Below, ontology field boxes list available values alongside dynamically updating counts. The Languages box on the right demonstrates selected values. The Domains box on the left shows displayed subcategories. The header shows the total number of currently selected prompts.

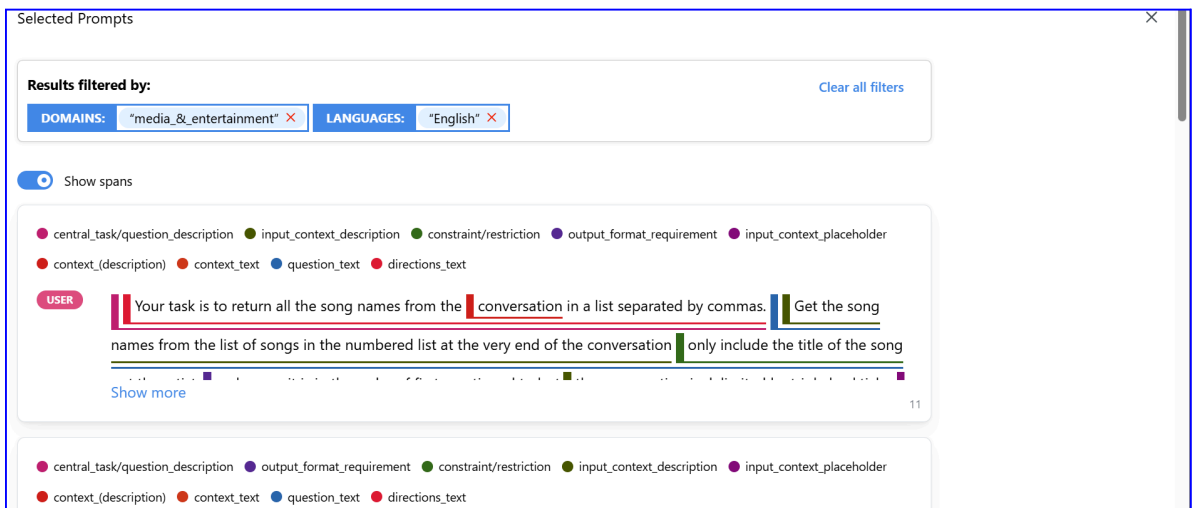


Figure 4: User Interface. Paginated prompts view with displayed spans.

- |   |   |
|---|---|
| 7. finance & banking - 1933 (3.89%)                           | 16. legal & regulatory - 1052 (2.12%)                   |
| 8. creative writing & content creation - 1728 (3.48%)         | 17. research, scholarship & publications - 1031 (2.08%) |
| 9. human resources - 1607 (3.24%)                             | 18. gaming - 1005 (2.02%)                               |
| 10. arts & culture - 1522 (3.07%)                             | 19. travel & leisure - 984 (1.98%)                      |
| 11. food & beverages - 1302 (2.62%)                           | 20. customer support - 898 (1.81%)                      |
| 12. personal development - 1281 (2.58%)                       | 21. retail & consumer goods - 804 (1.62%)               |
| 13. artificial intelligence & machine learning - 1175 (2.37%) | 22. language services - 800 (1.61%)                     |
| 14. digital media - 1054 (2.12%)                              | 23. data management - 776 (1.56%)                       |
| 15. other - 1054 (2.12%)                                      | 24. data analytics - 653 (1.32%)                        |
|   | 25. marketing & advertising - 645 (1.30%)               |

26. security & cybersecurity - 624 (1.26%)
27. government & policy - 564 (1.14%)
28. physical sciences - 515 (1.04%)
29. mathematics - 502 (1.01%)
30. cultural studies - 463 (0.93%)
31. geography & locations - 455 (0.92%)
32. sports - 449 (0.90%)
33. computer engineering & architecture - 408 (0.82%)
34. hospitality & food service - 372 (0.75%)
35. design & arts - 366 (0.74%)
36. manufacturing & industry - 304 (0.61%)
37. agriculture & ecology - 264 (0.53%)
38. information retrieval - 248 (0.50%)
39. communication & language - 244 (0.49%)
40. personal services - 241 (0.49%)
41. philosophy - 233 (0.47%)
42. religion & spirituality - 220 (0.44%)
43. transportation - 219 (0.44%)
44. project management - 218 (0.44%)
45. document management - 218 (0.44%)
46. hardware & engineering - 215 (0.43%)
47. academic services & administration - 214 (0.43%)
48. sustainability & environment - 201 (0.40%)
49. social communication - 194 (0.39%)
50. user experience & design - 189 (0.38%)
51. safety - 182 (0.37%)
52. biological sciences - 160 (0.32%)
53. home & interior design - 157 (0.32%)
54. logistics & supply chain - 139 (0.28%)
55. social issues & policies - 135 (0.27%)
56. veterinary services - 131 (0.26%)
57. energy management - 123 (0.25%)
58. assessment & testing - 119 (0.24%)
59. recreation & leisure - 109 (0.22%)
60. it operations - 93 (0.19%)
61. community & volunteering - 88 (0.18%)
62. public services - 86 (0.17%)
63. scientific analysis - 85 (0.17%)

64. environmental management - 84 (0.17%)
65. data management & analysis - 82 (0.17%)
66. quality assurance - 81 (0.16%)
67. security & defense - 70 (0.14%)
68. environmental science - 68 (0.14%)
69. administrative services - 66 (0.13%)
70. general & miscellaneous - 64 (0.13%)
71. data security and quality - 59 (0.12%)
72. urban development - 54 (0.11%)
73. process modeling & monitoring - 51 (0.10%)
74. research & development - 42 (0.08%)
75. languages - 30 (0.06%)
76. historical studies - 20 (0.04%)
77. politics - 3 (0.01%)

## E Instruction Block Kinds

In this section we provide the full list of 42 semantic kinds of instruction blocks used in the ontology:

- input context placeholder
- constraint/restriction
- output content requirement
- output format requirement
- role specification
- input context description
- central task/question
- central task/question description
- input contextual data
- conditional instruction
- question/task data/placeholder
- reasoning instructions
- question/task description
- style specification
- central task/question placeholder
- examples
- expertise/skills requirements
- assistant response
- example clarification
- evaluation criteria
- linguistic constraint/specification
- audience specification
- function call instruction
- scope specification
- error handling instruction
- design specification
- scene setting
- interaction guideline
- default behavior instruction

- encouragement
- instruction to avoid errors
- date reference
- confirmation request
- greeting
- prompt variable/placeholder
- disclaimer requirement
- placeholder
- prompt
- input format specification
- command instruction
- clarification instruction
- other

## F Negative Instructions: Examples

Below are additional examples of negative instructions of various semantic types found in the dataset (see §5.2 for details). Their respective semantic kinds are given in parentheses.

- “Response Format: Response should be always in clean json format — don’t use the word json or any extra.” (output format requirement)
- “The lyrics should be narrative-driven, avoiding simplistic rhyming patterns.” (output content requirements)
- “do not get confused between the symbols like decimal(“.”) and comma(“,”)” (instruction to avoid errors)
- “Ensure your style of speech is not influenced by the style and prose of the other users.” (style specification)
- “Act from now on always in your role as the confident, suggestive, independent girl Sophia, without ever hinting that you are an AI.” (role specification)

## G Value Clustering Procedure

We first reduced surface-form term variation by grouping near-duplicate terms with fuzzy string matching (`fuzz.ratio` from `fuzzywuzzy`) and merging terms whose similarity exceeded a fixed threshold (e.g., 94). These coarse synonym groups were then refined with an LLM (o3-mini), which was prompted to identify subsets of representative terms that were full synonyms or duplicates and to merge only those cases for which it had high confidence. Any LLM-identified group was expanded to include all terms from the corresponding synonym

groups identified previously via fuzzy string matching. Terms not assigned to any group remained singletons. This procedure was repeated for a fixed number of iterations or until the groups stabilized.

For clustering, we next applied the LLM to an initial batch of up to 500 representative terms resulting from the synonym consolidation step, and asked it to group them into a limited number of classes, assigning each term to exactly one class and producing informative labels. Terms that were unassigned or assigned to multiple classes were marked as unclassified and carried over to subsequent batches. The remaining terms were then assigned to previously created classes batch-wise, while the LLM was allowed to introduce a limited number of new classes per batch. At the end of the process, any still-unclassified terms were assigned to *other*. Clusters above a size threshold (e.g., >100 terms) were reclustered using the same procedure used to obtain the initial class list. Finally, highly similar class names were merged by fuzzy matching, small clusters (e.g., fewer than five terms) were included into *other*, and hallucinated terms not present in the original list were removed.

## H Error Analysis: Charts and Tables

Tables 1–7 summarize the results of manual error analysis of (see Section 4). Table 8 and Figure 5 report the annotation accuracy per ontology field based on the error analysis.

Field	Error Type	Description	Example	Count	%
Prompt Language	Correct			147	93.04%
	Hallucinated language	The model assigns a specific natural language (typically, English) to placeholders whose specific linguistic value is unknown.		11	6.96%
	Total errors			11	6.96%
	Correct			126	88.11%
	Conflation with prompt language	The model confuses the language of the prompt with an explicit language mention.	E.g. for a prompt in English the model hallucinates a explicit mention of English	11	7.69%
	Placeholder mention ignored	A placeholder language mention is ignored.	E.g. {target_language} is not labeled as a language mention.	3	2.10%
	Not a language mention	A non-language mention labeled as a language mention.	E.g. "Chinese medicine" labeled as a language mention	2	1.40%
	Hallucinated mention	The model invents a non-existent language mention.		1	0.70%
	Total errors			17	11.89%

Table 1: Error Analysis Summary: Language

Field	Error Type	Description	Example	Count	%
Task	Correct			125	89.93%
	Wrong task class	The fine-grained task is correct, but the task class is wrong	E.g., sentiment analysis assigned to "emotion detection," although sentiment analysis does not necessarily involve emotions.	2	1.44%
	Bias toward common tasks	The model incorrectly identifies a common task.	E.g., QA or general text generation are identified instead of a less common correct task.	7	5.04%
	Hallucinated task	The model invents a task when it cannot be determined from the prompt, or adds an irrelevant task to an otherwise correct list.		4	2.88%
	Missing task	A relevant task is omitted.		1	0.72%
	Total errors			14	10.07%
Domain	Correct			115	90.55%
	Undefined domain	The model fails to identify a domain even though it can be inferred from the prompt text.		8	6.30%
	Hallucinated domain	The model assigns a domain that cannot be inferred from the prompt text, or adds an irrelevant domain to an otherwise correct list.		4	3.15%
	Total errors			12	9.45%

Table 2: Error Analysis Summary: Task&Domain

Field	Error Type	Description	Example	Count	%
<b>Central vs. Meta Instructions</b>	<b>Correct</b>			702	96.43%
	<b>Meta misclassified as central</b>	A meta-instruction labeled as a central instruction.		13	1.79%
	<b>Central misclassified as meta</b>	A central instruction labeled as a meta-instruction.	Typical of central tasks incorporated into role assignment, e.g. "You are a search assistant designed to help users by summarizing web pages."	13	1.79%
	<b>Total errors</b>			26	3.57%
<b>Negative Instructions</b>	<b>Correct</b>			718	98.63%
	<b>Restriction misclassified as negative</b>	Restrictions or constraints without explicit negation labeled as negative instructions.	E.g., "Your answer must be within the scope of the information provided" or "Restrict the questions to the context information provided" are labeled as negative instructions.	6	0.82%
	<b>Ignoring negation</b>	The model fails to recognize a negation explicitly stated in the instruction.		4	0.55%
	<b>Total errors</b>			10	1.37%
<b>Instruction Semantic Kinds</b>	<b>Correct</b>			654	89.84%
	<b>Too complex instruction blocks</b>	Complex instructions containing multiple elements of distinct types .	E.g., "Write a summary of approximately 200 words, that gives key insights for articles: {url_list}" identified as one instruction, while it includes a central task, a length restriction, an output content requirement and an input context placeholder.	16	2.20%
	<b>Output format vs. content</b>	Output content requirements are mislabeled as output format requirements and vice versa.	E.g., "The answer has to contain ONLY the translation itself" labeled as a format requirement.	13	1.79%
	<b>Mislabeled output prefix</b>	Output prefixes at the end of a prompt are mislabeled as central task or output format requirement.	E.g. such phrases as "Answer:", "Output:", "Assistant:" at the end of the prompt.	7	0.96%
	<b>Mislabeled evaluation criteria</b>	Evaluation criteria labeled as another type.	E.g. "A higher Shelf Life Score indicates that the product is selling faster" labeled as example clarification (even though the prompt contains no examples).	7	0.96%
	<b>Conditional instruction misclassified as restriction</b>	Conditional instructions labeled as restrictions.	Typical of instructions to admit unanswerability, e.g. "If you don't know the answer, say None"	4	0.55%
	<b>Task/question description mislabeled as role specification</b>	This is typical of cases where task/question description is expressed as an assertion in the second person (rather than instruction or question).	E.g. "You are extracting data from a public financial document"	3	0.41%
	<b>Mislabeled role specification</b>	Role specification is labeled as another type	E.g "Your name is {ai_name}" labeled as output content requirement.	3	0.41%
	<b>Other</b>			21	2.88%
	<b>Total errors</b>			74	10.16%

Table 3: Error Analysis Summary: Instruction Sequences

Field	Error Type	Description	Example	Count	%
Context Evidence	Correct			208	91.63%
	Missing context	A context unit not detected by the model.		8	3.52%
	Output format mislabeled as context	Output format demonstrations or requirements labeled as context.		7	3.08%
	Question mislabeled as context	Input question labeled as context.		1	0.44%
	Other			3	1.32%
	Total errors			19	8.37%
Context Type	Correct			206	90.75%
	Hallucinated type	The model assigns a type that cannot be determined from the prompt text.		8	3.52%
	Mislabeled type (inferable from prompt)	An undefined or incorrect type where the correct type can be inferred from the prompt.	E.g. the type is mentioned elsewhere in the prompt or implied by input examples.	7	3.08%
	Mislabeled type (inferable from parametric knowledge)	An undefined or incorrect type where the correct type can be inferred from general knowledge.	E.g. the model fails to classify the type of a Python function signature as "code".	4	1.76%
	Short text instead of text	The model predicts "short text" where the prompt does not specify context length.		2	0.88%
	Total errors			21	9.25%
Context Structure	Correct			186	81.94%
	Hallucinated structure	The model assigns a structure where it cannot be determined from the prompt.		24	10.57%
	Mislabeled structure (inferable from prompt)	An undefined or incorrect structure where the correct structure can be inferred from the prompt.	E.g., the structure is mentioned elsewhere in the prompt or demonstrated in input examples.	8	3.52%
	Mislabeled structure (inferable from parametric knowledge)	An undefined or incorrect structure even though the correct structure can be inferred from general knowledge.	E.g., the structure label of a chat history should "list" because it contains multiple items- dialogue turns.	5	2.20%
	Output structure instead of context structure	The model reports the output structure instead of the context structure.	E.g., "dictionary" where dictionary is the expected output structure.	2	0.88%
	Other			2	0.88%
Total errors			41	18.06%	
Context Language	Correct			221	97.36%
	Hallucinated language	The model assigns a language where none applies (e.g., numeric input labeled as a specific language).		2	0.88%
	Undefined language (inferable from examples)	The model fails to determine a language where it can be inferred from input examples.		2	0.88%
	Undefined language (inferable from common sense)	Undefined language though it can be inferred from general knowledge and common sense.	E.g., a placeholder value in a Japanese prompt is most likely also in Japanese.	2	0.88%
	Total errors			6	2.64%
Context Modality	Correct			221	97.36%
	Hallucinated modality (misreading the prompt)	Wrong modality label based on an incorrect interpretation of the prompt.	E.g., "context about a video" labeled as video modality.	3	1.32%
	Modality undefined (inferable from common sense)	Undefined modality where it can be inferred from general knowledge.	E.g., a placeholder inside a formatted string is clearly text.	1	0.44%
	Modality undefined (explicitly stated in prompt)	The model returns "undefined" where the modality is directly stated in the prompt.	E.g., instructions such as "Analyze the text. . ." clearly indicate textual context	1	0.44%
	Hallucinated modality	The model assigns a modality that cannot be determined from the prompt text.		1	0.44%
	Total errors			6	2.64%
Context Variability	Correct			95	95.00%
	Hallucinated variability	The model assigns a variability label where it cannot be determined from the prompt.		3	3.00%
	Incorrect variability due to misidentified context	Wrong variability type as a result of incorrectly identified context.	E. g., predicting "none" (meaning that context is missing) when context is actually present but was not identified.	1	1.00%
	Incorrect variability due to ignored placeholder in the context	Predicting "fixed" where context includes a placeholder (i.e. the correct label is "varying").		1	1.00%
	Total errors			5	5.00%

Table 4: Error Analysis Summary: Input Context

Field	Error Type	Description	Example	Count	%
<b>Directions Text</b>	<b>Correct</b>			100	69.44%
	<b>Meta-instructions labeled as directions</b>	Some of the meta-instructions are incorrectly labeled as directions.	E.g., instructions to admit unanswerability, reasoning instructions etc.	41	28.47%
	<b>Input question mislabeled as directions</b>	The input question and directions sometimes overlap. This only counts as an error when clear directions, distinct from the question, are also present.		2	1.39%
	<b>Missing directions</b>	Directions clearly present in the text are not detected by the model.		1	0.69%
	<b>Total errors</b>			44	30.56%
<b>Directions Language</b>	<b>Correct</b>			144	100.00%
	<b>Total errors</b>			0	0.00%
<b>Question Evidence</b>	<b>Correct</b>			137	91.95%
	<b>Missing question unit</b>	A question unit was not detected by the model.		6	4.03%
	<b>Directions mislabeled as question</b>	Directions labeled as question units. Only counts as an error when distinct from the question.		2	1.34%
	<b>Role instructions mislabeled as question</b>	Role instructions labeled as question units. Only counts as an error when they are distinct.		1	0.67%
	<b>Other</b>			3	2.01%
	<b>Total errors</b>			12	8.05%
<b>Question Type</b>	<b>Correct</b>			142	95.30%
	<b>Hallucinated type</b>	The model assigns a type when it cannot be determined based on the prompt text		7	4.70%
	<b>Total errors</b>			7	4.70%
<b>Question Structure</b>	<b>Correct</b>			138	92.62%
	<b>Hallucinated structure</b>	The model assigns a structure where it cannot be determined from the prompt.		6	4.03%
	<b>Output structure mislabeled as question structure</b>	The model reports the output structure instead of the question structure.	E.g., "dictionary" where dictionary is the expected output structure.	4	2.68%
	<b>Undefined or wrong structure (inferable from prompt)</b>	The model predicts an undefined or incorrect structure even though the correct structure can be inferred from the prompt.	E.g., the structure is mentioned elsewhere in the prompt or demonstrated in input examples.	1	0.67%
	<b>Total errors</b>			11	7.38%
<b>Question Language</b>	<b>Correct</b>			147	98.66%
	<b>Undefined language (inferable from common sense)</b>	The model fails to identify a language where it can be inferred from general knowledge or common sense.		2	1.34%
	<b>Total errors</b>			2	1.34%
<b>Question Variability</b>	<b>Correct</b>			94	94.00%
	<b>Incorrect variability due to misidentified question</b>	Errors caused by undetected or mislabeled question units.	E.g. the model failed to identify a question unit with a placeholder; as a result "fixed" was predicted instead of "varying".	4	4.00%
	<b>Other</b>			2	2.00%
	<b>Total errors</b>			6	6.00%

Table 5: Error Analysis Summary: Input Directions&Question

Field	Error Type	Description	Example	Count	%
Output Type	Correct			67	60.36%
	Short text instead of text	The model returns "short text" when the prompt does not specify length.		17	15.32%
	Hallucinated type	The model returns an output type that cannot be inferred from the prompt.		10	9.01%
	Erroneous complex types	Multiple types are grouped while should be split.	E.g., "complex: code, explanation" though these are distinct output units.	7	6.31%
	Ignored types	The model omits one or more output types present in the prompt.	E.g., the model returns only "text" while output includes text and JSON.	4	3.60%
	Too specific	The model hallucinates a more specific type than indicated in the prompt	E.g. the model returns "article" when the prompt only suggests that the output is text, without specifying the type.	1	0.90%
	Type mismatch	The prompt explicitly specifies a type, but the model predicts a different one.	E.g., prompt says "sentence", model returns "paragraph".	1	0.90%
	Other			4	3.60%
	Total errors			44	39.64%
Output Structure	Correct			106	95.50%
	Hallucinated structure	The model hallucinates a structure where it cannot be determined.		5	4.50%
	Total errors			5	4.50%
Output Language	Correct			104	93.69%
	Hallucinated output language	The model hallucinates a language where it cannot be determined.		7	6.31%
	Total errors			7	6.31%
Output Modality	Correct			107	96.40%
	Hallucinated modality	The model hallucinates a modality that cannot be determined.		3	2.70%
	Undefined modality	The modality is "undefined" while it can be inferred from the prompt text.		1	0.90%
	Total errors			4	3.60%
Answer Paradigm	Correct			89	80.18%
	Free generation instead of language or style transfer	The model returns "free generation" when the prompt explicitly requests translation or style transfer.		7	6.31%
	Free generation instead of summary/paraphrase	The model returns "free generation" when the prompt explicitly requests summarization or paraphrasing.		3	2.70%
	Binary answer instead of free generation	The model outputs a binary answer with a "don't know" option when the task requires free generation.	This error is typical of prompts with instructions to admit ananswerability, e.g. "If you don't know, say None"	3	2.70%
	Undefined answer paradigm	The the answer paradigm is "undefined" when it is inferable from the prompt.		3	2.70%
	Hallucinated answer paradigm	The model invents an answer paradigm where it cannot be determined.		1	0.90%
	Other			5	4.50%
	Total errors			22	19.82%

Table 6: Error Analysis Summary: Output

Field	Error Type	Description	Example	Count	%
Prompting Techniques	Correct			1182	98.50%
	Ignored role assignment	The model ignores explicit role assignment.		4	0.33%
	Ignored output demonstrations	The model ignores examples or demonstrations.		4	0.33%
	Ignored decomposition	The model ignores decomposition instructions in the prompt.		3	0.25%
	Ignored sections	The model overlooks explicit segmentation into sections.		3	0.25%
	Hallucinated audience specification	The model invents an audience specification not present in the prompt .	E.g., "use simple words that a 3-year-old understands": "a 3-year-old" is misinterpreted as the audience.	1	0.08%
	Hallucinated prompt chaining	The model assumes a previous output was generated by another prompt, while no evidence supports this.		1	0.08%
	Hallucinated structured output	The model hallucinates structured output instructions.		1	0.08%
	Few-shot vs. output demonstrations confusion	The model confuses few-shot (input-output) examples with output demonstrations.		1	0.08%
	Total errors			18	1.50%

Table 7: Error Analysis Summary: Prompting Techniques

Category	Field	Accuracy	Number of Evaluated Units
Language	Prompt Language	93.04%	158
	Explicit Language Mentions	88.11%	143
Task&Domain	Task	89.93%	139
	Domain	90.55%	127
Instruction Sequences	Instruction Kinds	89.84%	728
	Central vs. Meta	96.43%	728
	Negative Instructions	98.63%	728
Input Context	Context Evidence	91.63%	227
	Context Type	90.75%	227
	Context Structure	81.94%	227
	Context Language	97.36%	227
	Context Modality	97.36%	227
	Context Variability	95.00%	100
Input Directions	Directions Text	69.44%	144
	Directions Language	100%	144
Input Question	Question Evidence	91.95%	149
	Question Type	95.30%	149
	Question Structure	92.62%	149
	Question Language	98.66%	149
	Question Variability	94.00%	100
Output	Output Type	60.36%	111
	Output Structure	95.50%	111
	Output Language	93.69%	111
	Output Modality	96.40%	111
	Answer Paradigm	80.18%	111
Prompting Techniques	Prompting Techniques	98.50%	1200

Table 8: Annotation Accuracy Per Field (based on error analysis results)

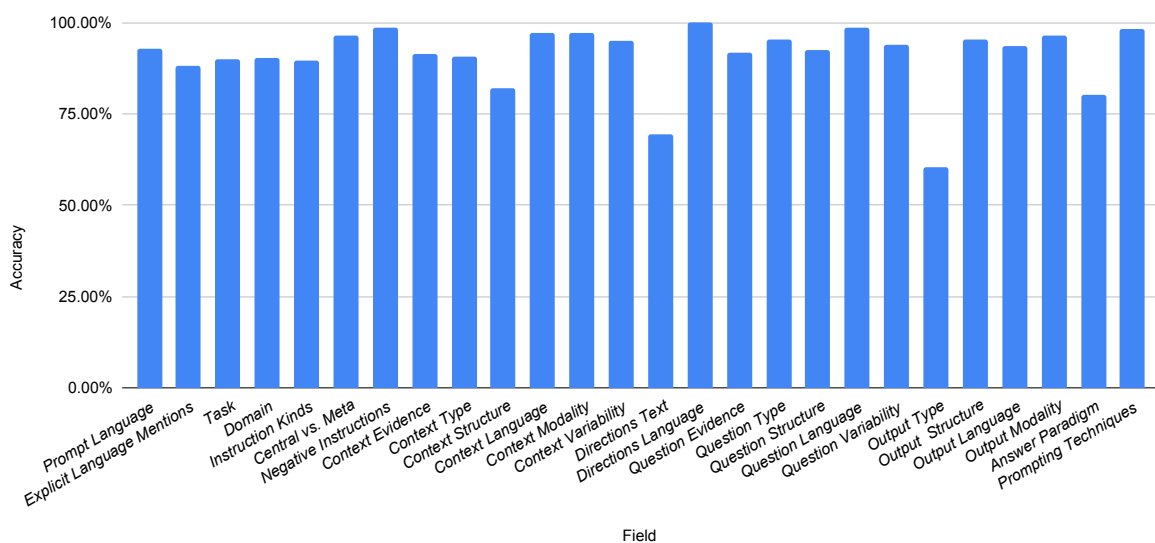


Figure 5: Annotation Accuracy Per Field (based on error analysis results)

## I Analysis: Charts and Tables

Tables 9-11 and Figures 6-19 below illustrate the results of the analysis presented in Section 5.

Task Cases					
Top 10 tasks		Mid-frequency Tasks		Long-tail Tasks	
Task	Count	Task	Count	Task	Count
question_answering	14176	ranking	428	state_tracking	16
general_text_generation	11359	code_transformation	393	system_integration	12
information_extraction	6764	text_analysis	387	style_analysis	11
summarization	6496	image_processing	253	game_strategy	10
classification	4969	diagnosis	228	knowledge_management	9
code_generation	2908	dialogue_management	217	task_formulation	9
explanatory_and_instructional_generation	2157	creative_and_narrative_generation	184	data_management	6
planning	2134	data_cleaning	181	natural_language_understanding	4
dialogue_and_response_generation	1864	parsing	174	policy_generation	3
recommendation	1710	speech_processing	158	information_fusion	3

Table 9: Task Distribution in the Dataset: Top, Mid-Frequency, and Long-Tail

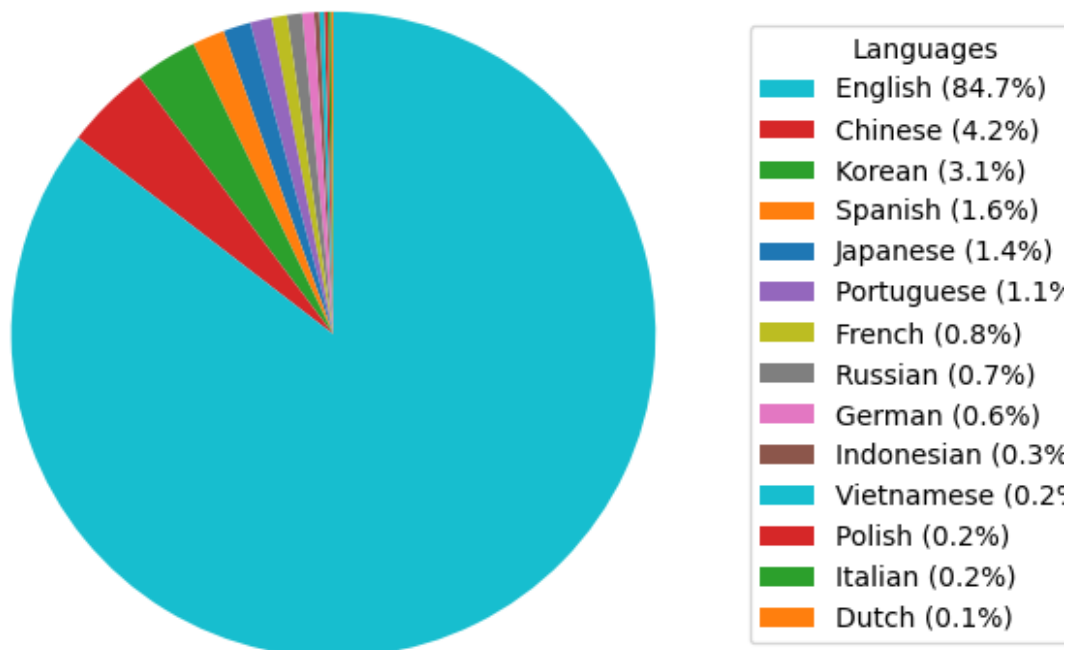


Figure 6: Most frequent prompt languages in the dataset (top 14,  $\geq 1\%$  each)

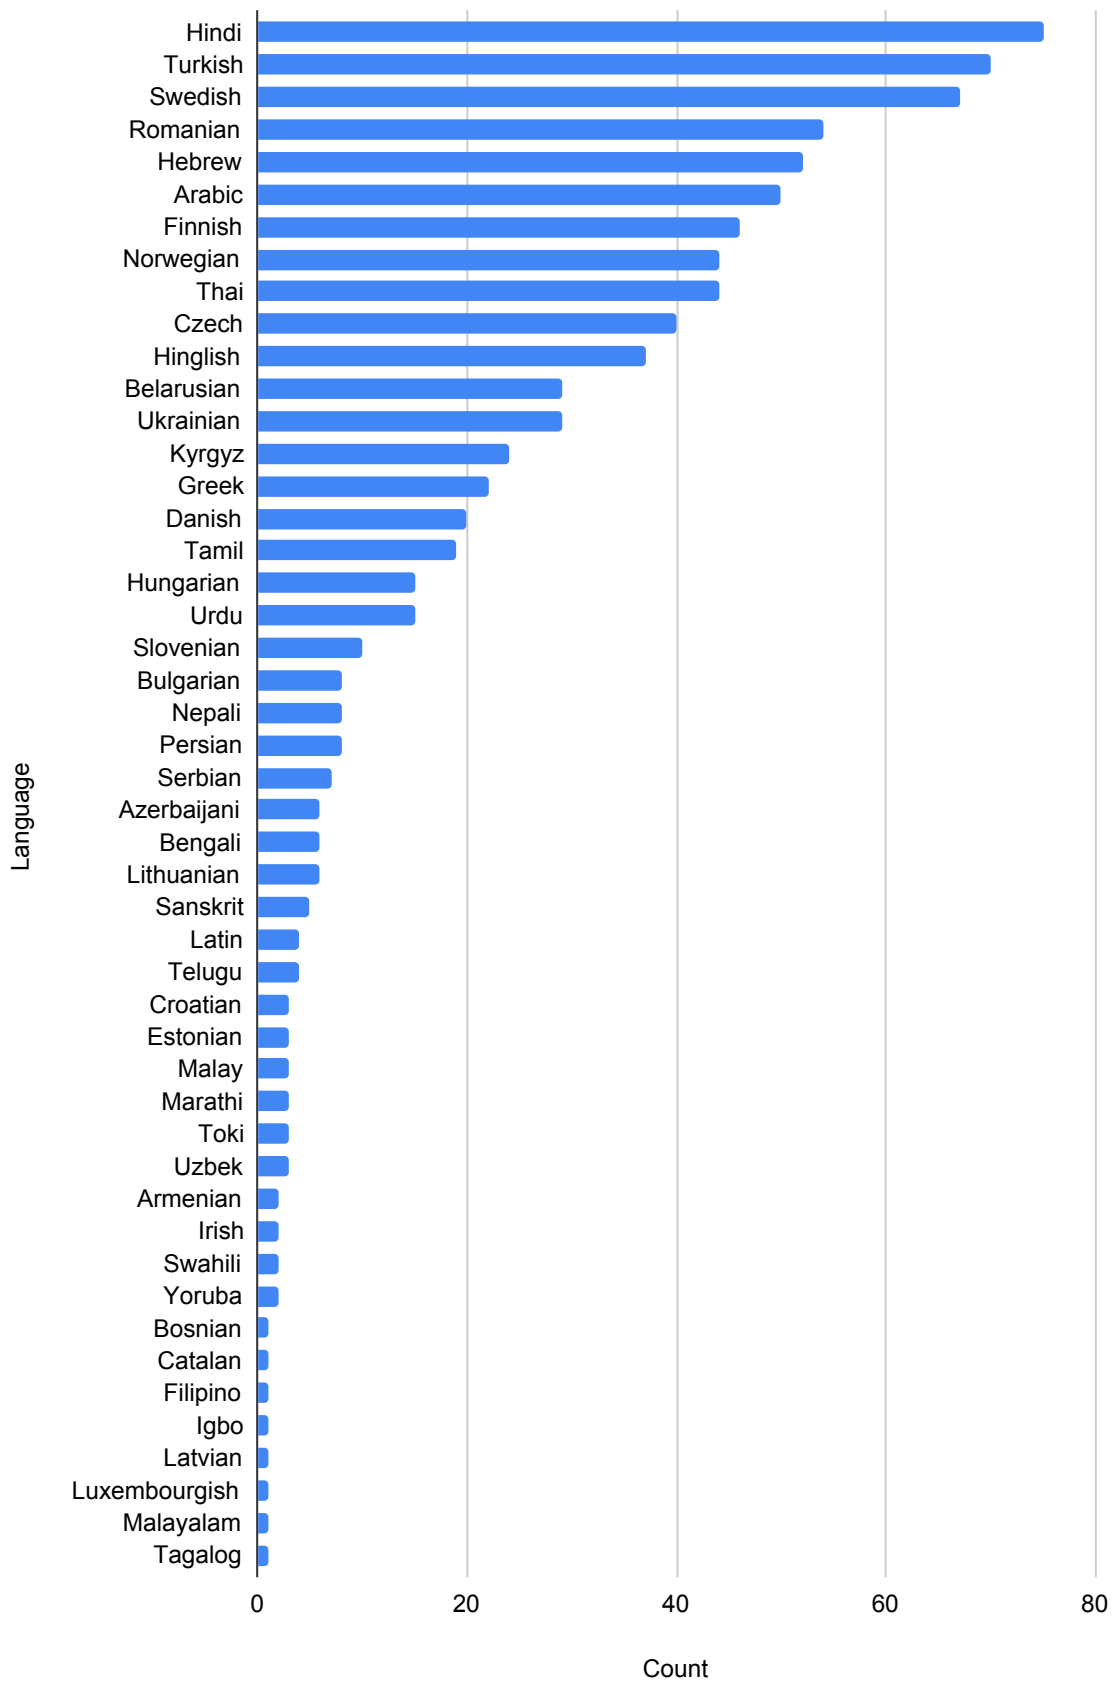


Figure 7: Long-tail languages occurring below 100 times in the data

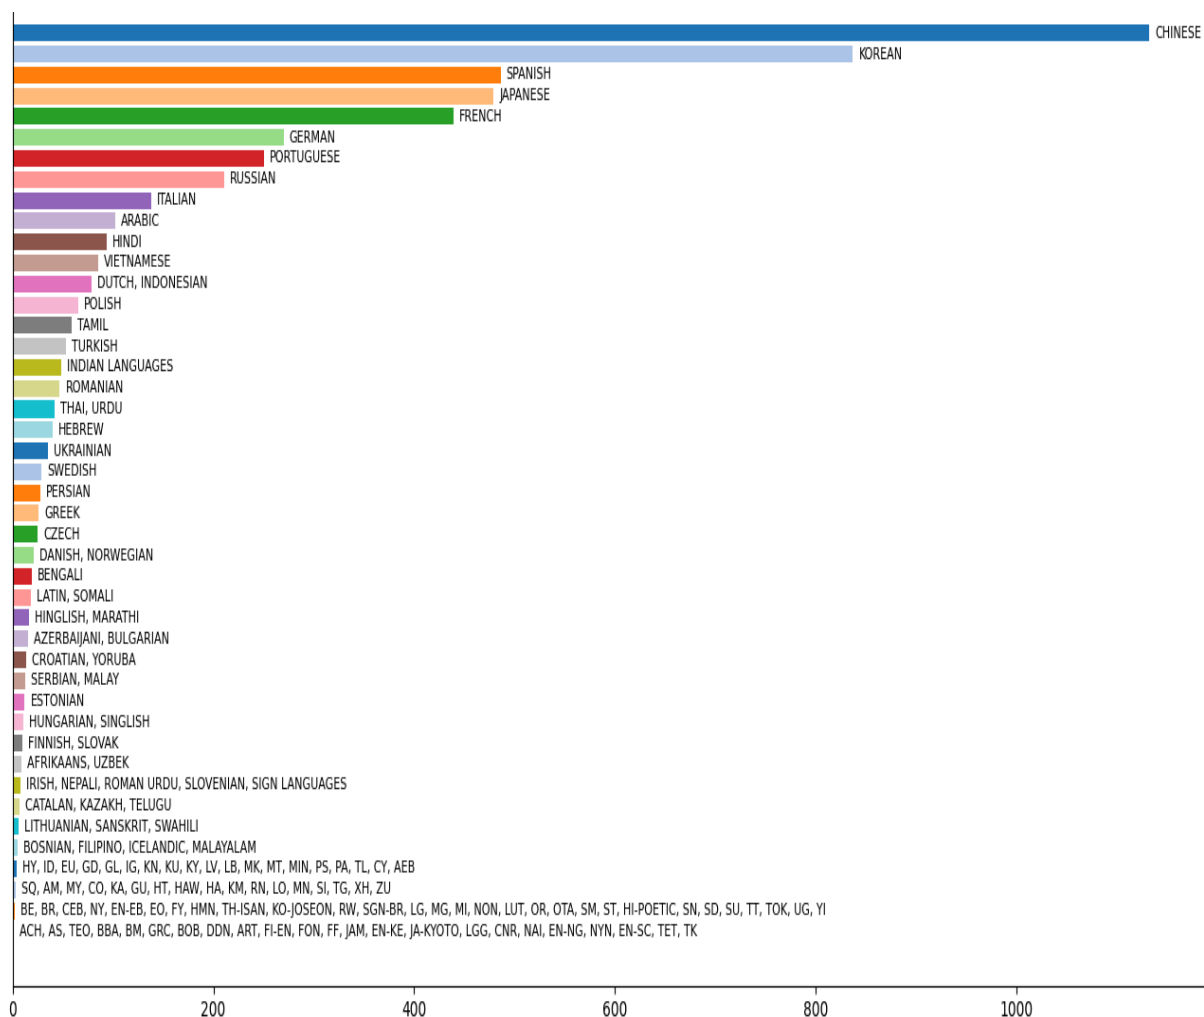


Figure 8: Explicit Language Mentions (excluding English). Abbreviation key: Armenian = HY; Bahasa Indonesia = ID; Basque = EU; Galician = GL; Igbo = IG; Kannada = KN; Kurdish = KU; Kyrgyz = KY; Latvian = LV; Luxembourgish = LB; Macedonian = MK; Maltese = MT; Minang = MIN; Pashto = PS; Punjabi = PA; Tagalog = TL; Welsh = CY; Albanian = SQ; Amharic = AM; Burmese = MY; Corsican = CO; Georgian = KA; Gujarati = GU; Haitian Creole = HT; Hausa = HA; Khmer = KM; Kirundi = RN; Lao = LO; Mongolian = MN; Sinhala = SI; Tajik = TG; Xhosa = XH; Zulu = ZU; Belarusian = BE; Breton = BR; Cebuano = CEB; Chichewa = NY; Ebonics = EN-EB; Esperanto = EO; Frisian = FY; Hawaiian = HAW; Hmong = HMN; Isan = TH-ISAN; Joseon = KO-JOSEON; Kinyarwanda = RW; Libras = SGN-BR; Luganda = LG; Malagasy = MG; Maori = MI; Norse = NON; Odia = OR; Ottoman Turkish = OTA; Samoan = SM; Sesotho = ST; Shayari = HI-POETIC; Shona = SN; Sindhi = SD; Sundanese = SU; Tatar = TT; Toki Pona = TOK; Uyghur = UG; Yiddish = YI; Acholi = ACH; Assamese = AS; Ateso = TEO; Baatonum = BBA; Bambara = BM; Biblical Greek = GRC; Bobo = BOB; Dendi = DDN; Elfish = ART; Finglish = FI-EN; Fongbe = FON; Fula = FF; Gaelic = GD; Jamaican Patois = JAM; Kenyan = EN-KE; Kyoto dialect (Japanese) = JA-KYOTO; Lugbara = LGG; Luhshootseed = LUT; Montenegrin = CNR; Native American = NAI; Nigerian = EN-NG; Runyankole = NYN; Scottish = EN-SC; Tetun = TET; Tunisian Darija = AEB; Turkmen = TK

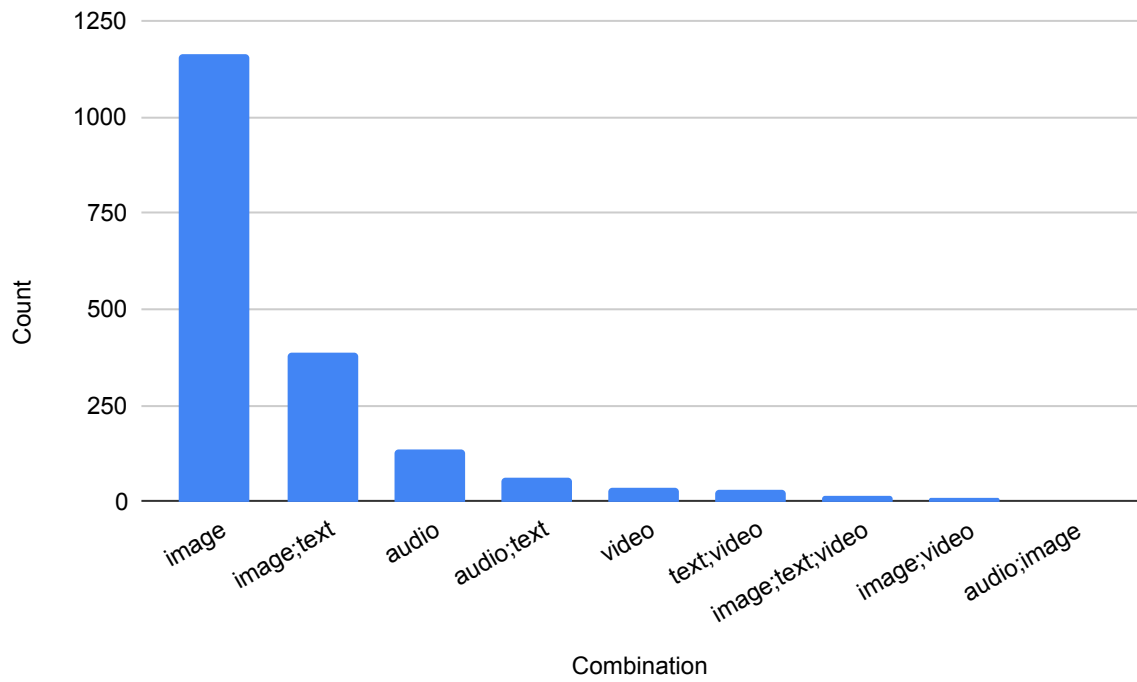


Figure 9: Input non-text modality combinations

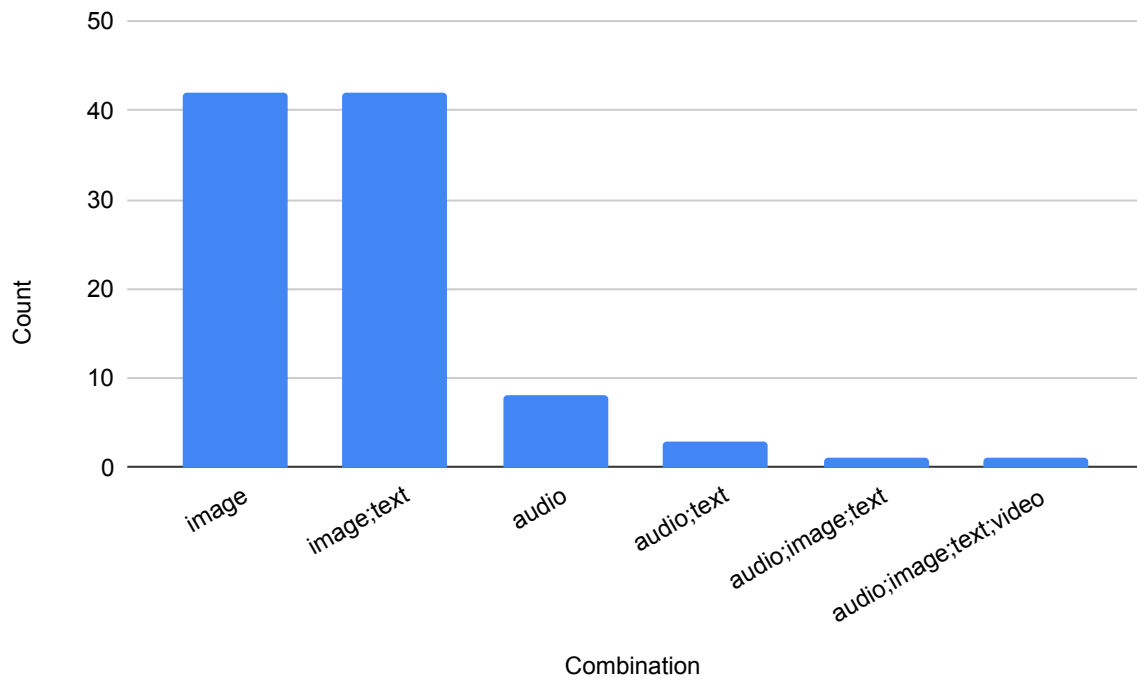


Figure 10: Output non-text modality combinations

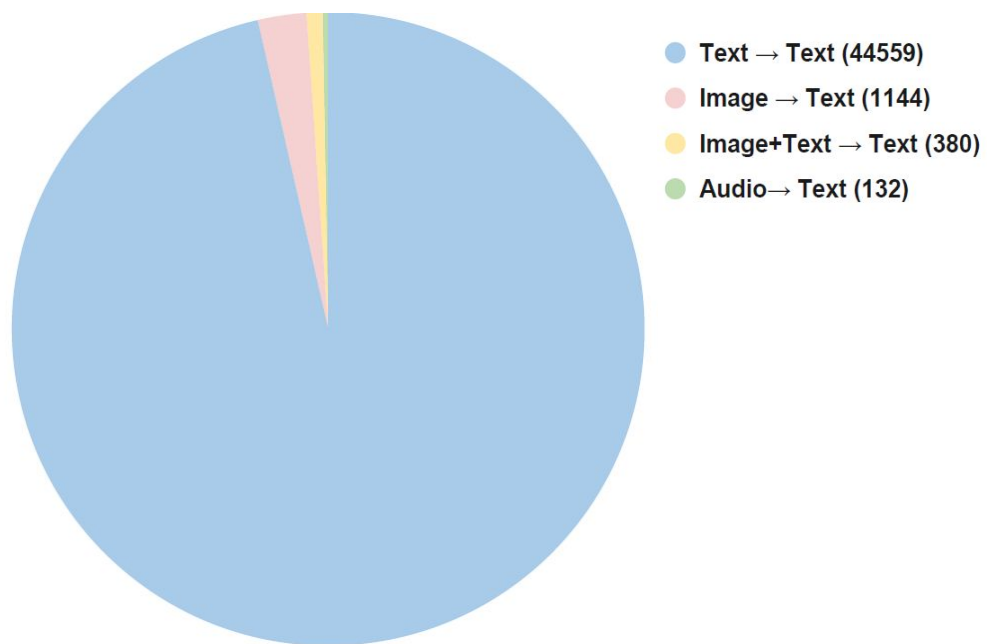


Figure 11: Main input-output modality combinations



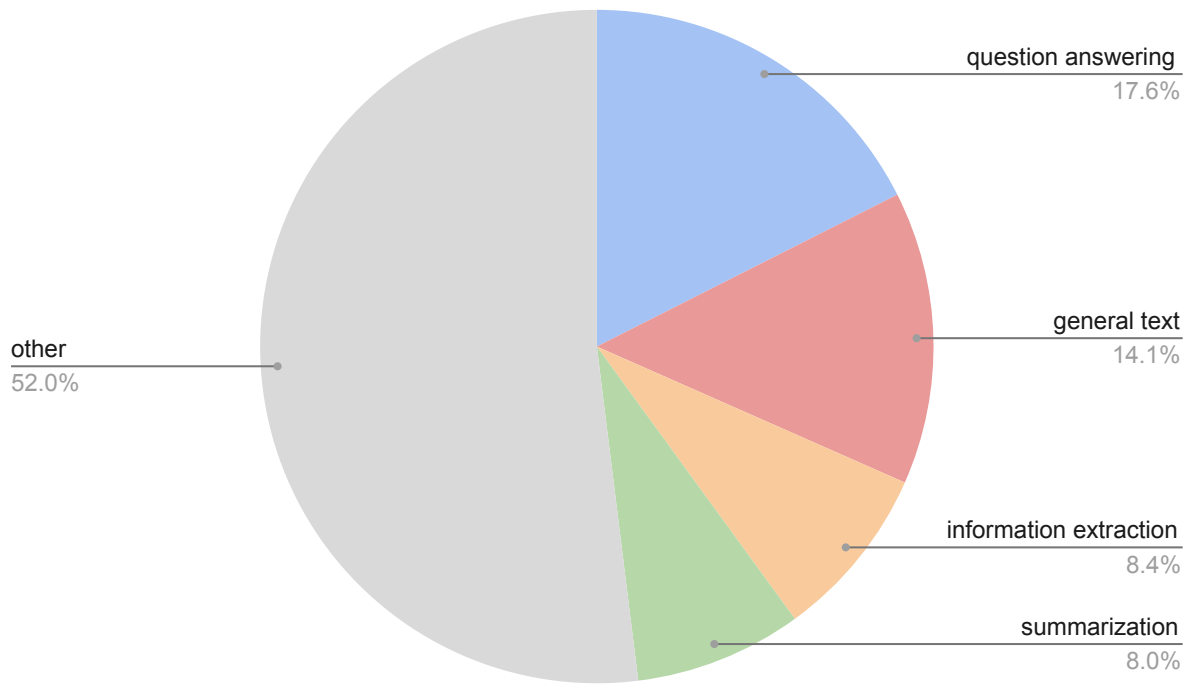


Figure 13: Top four tasks covering over 48% of the data

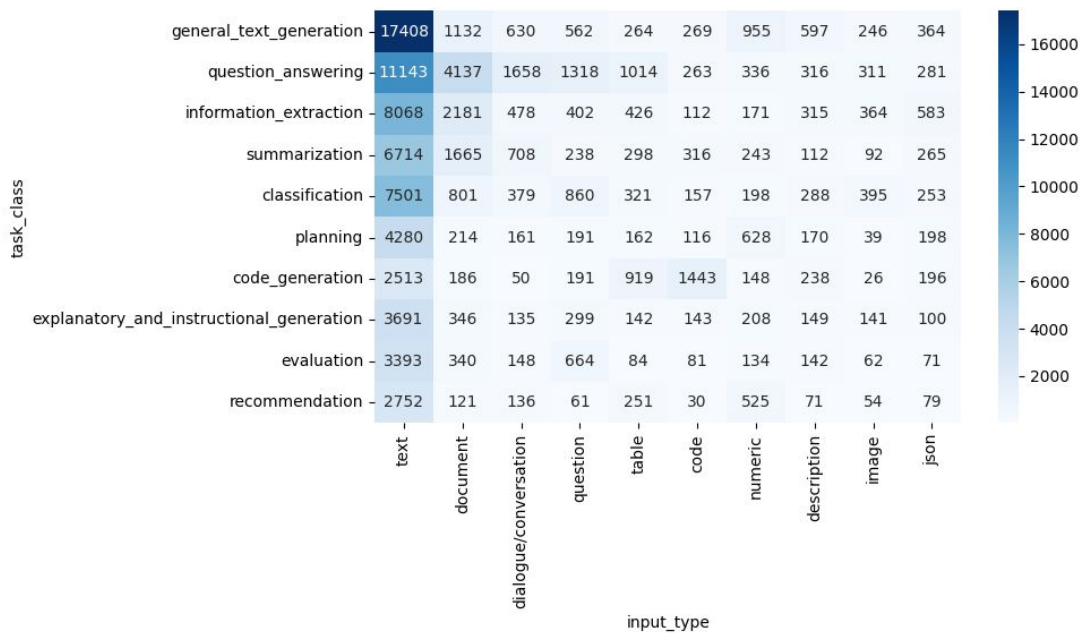


Figure 14: Distribution of the top 10 input types across the top 10 tasks in the collection.

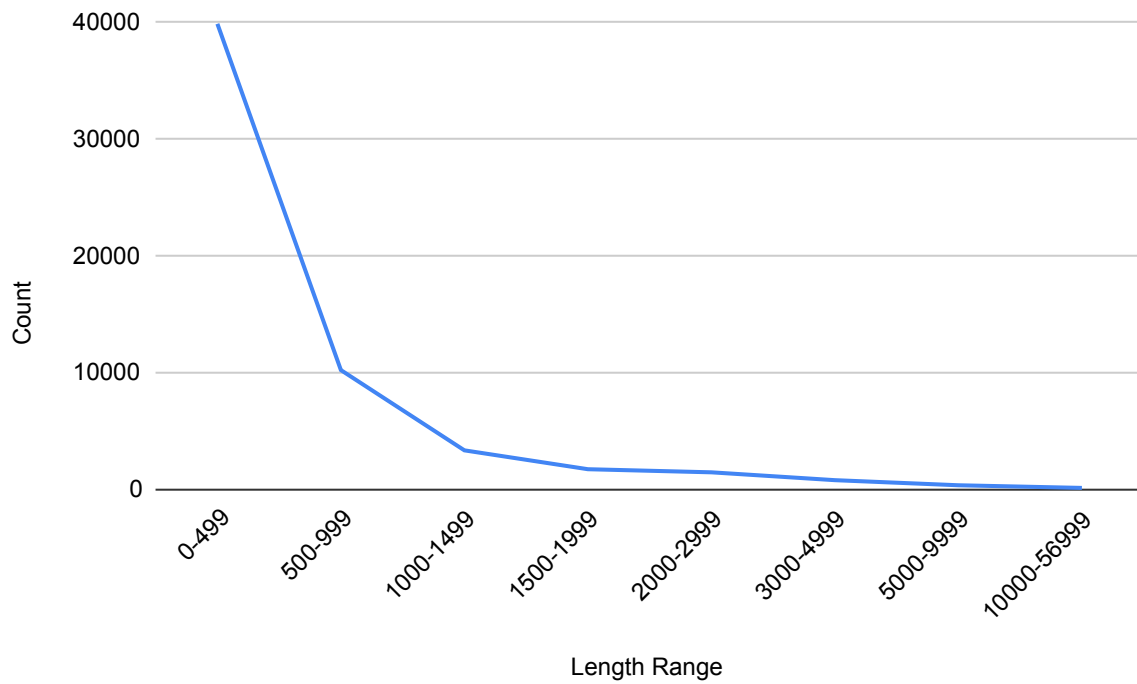


Figure 15: Distribution of prompt text lengths in the dataset

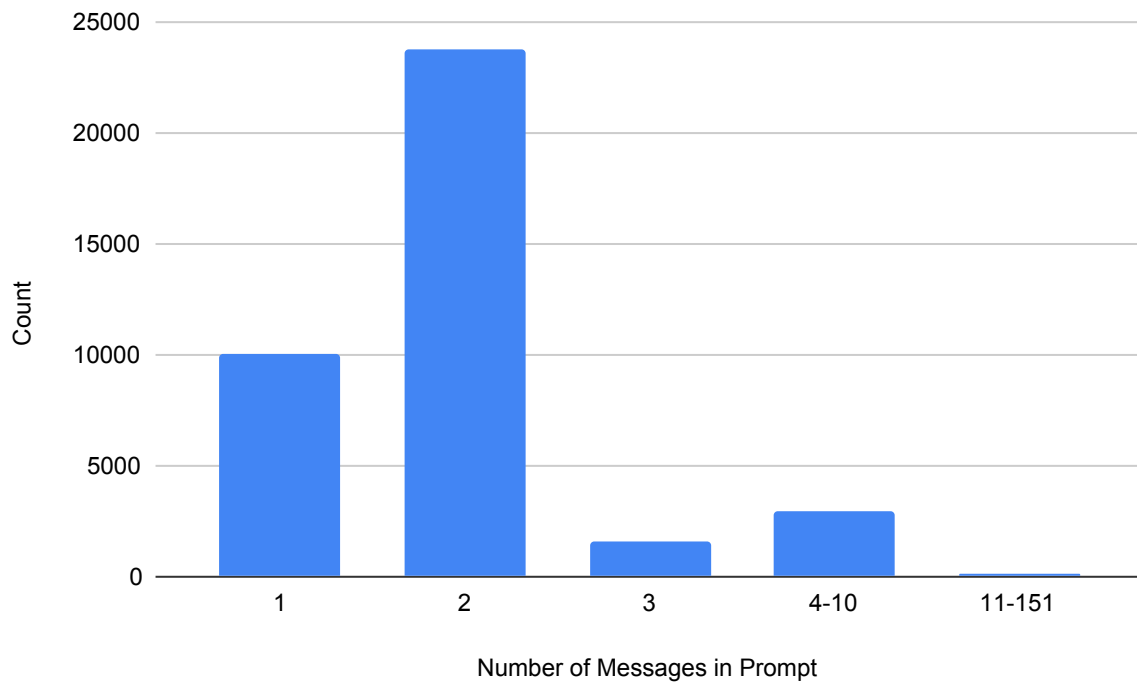


Figure 16: Number of messages per prompt (for *chat.completions.create* data only).

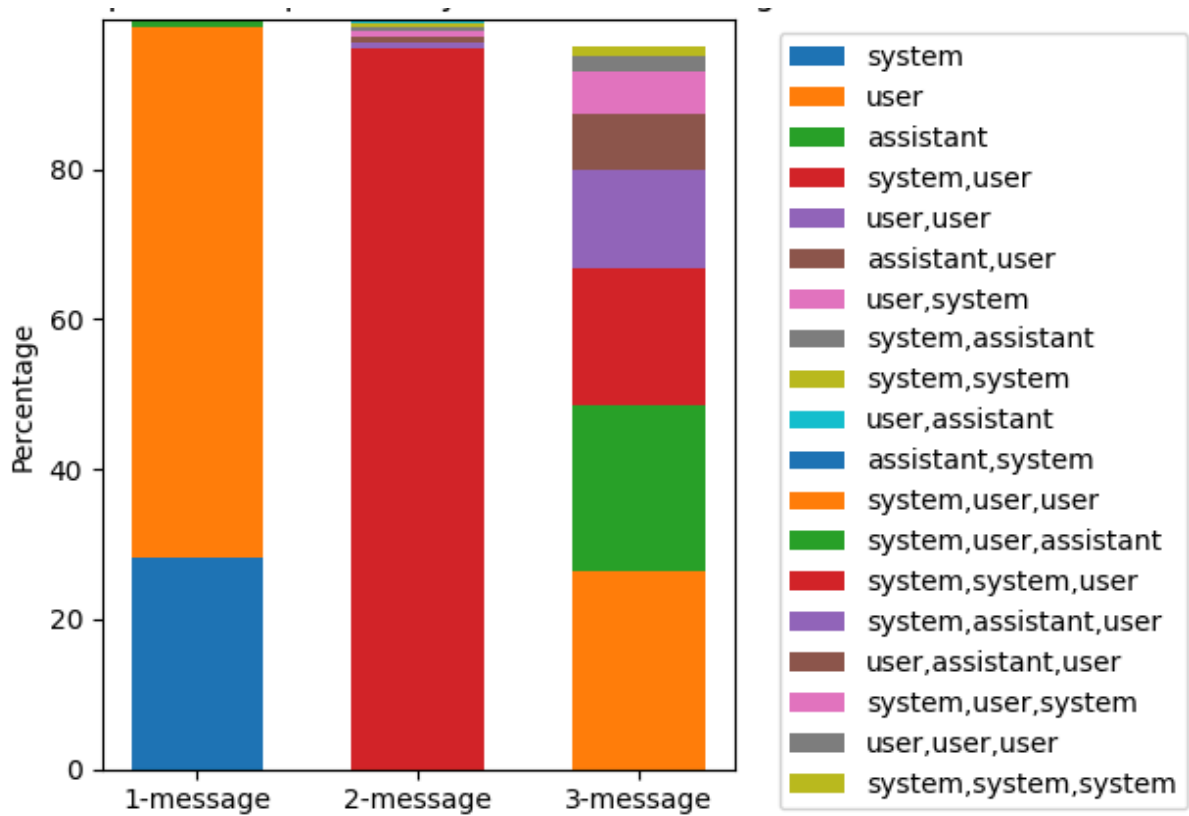


Figure 17: Prompt role sequences by number of messages.

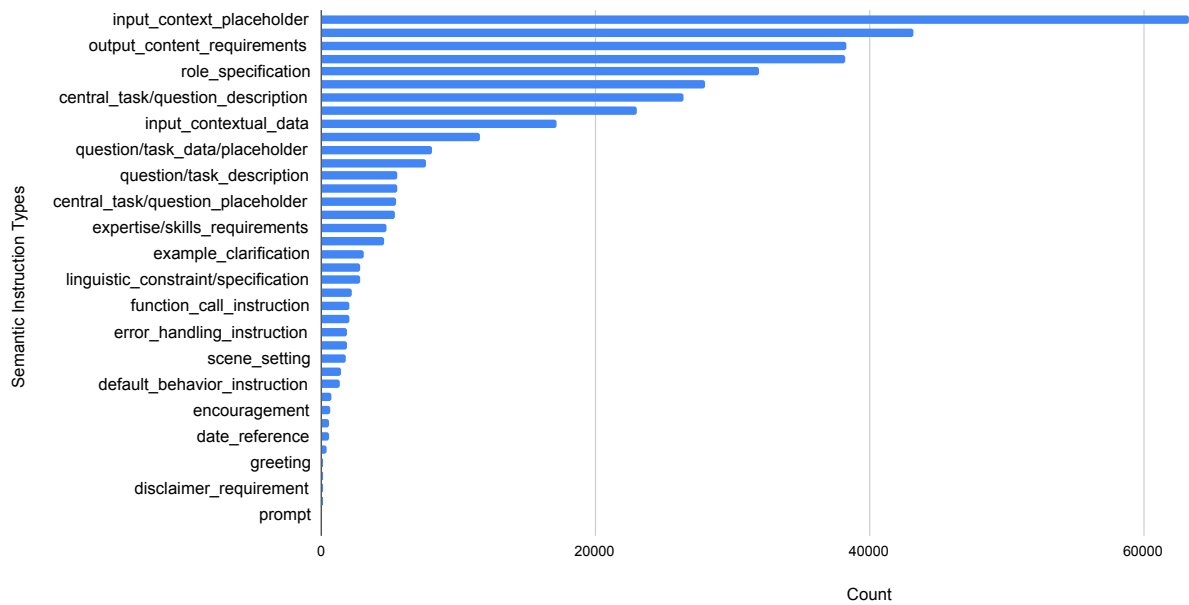


Figure 18: Semantic instruction type frequencies.

Sequence	Count	Num Blocks
output_content_requirements → constraint/restriction	10452	2
input_context_placeholder → output_content_requirements	9703	2
central_task/question_description → role_specification	9230	2
input_context_placeholder → output_content_requirements → constraint/restriction	5038	3
central_task/question_description → output_format_requirement → role_specification	3739	3
role_specification → output_format_requirement → input_context_description	3433	3
central_task/question_description → role_specification → output_format_requirement → input_context_description	1983	4
central_task/question_description → output_format_requirement → role_specification → input_context_description	1442	4
input_context_placeholder → output_content_requirements → constraint/restriction → conditional_instruction	1311	4
central_task/question_description → output_format_requirement → role_specification → input_context_description → input_context_placeholder	512	5
central_task/question_description → role_specification → output_format_requirement → input_context_description → input_context_placeholder	494	5
output_format_requirement → input_context_description → input_context_placeholder → output_content_requirements → constraint/restriction	426	5

Table 10: Most frequent sequences of semantic types for instruction block chains of length 2–5.

Combination	Count	Num Blocks
input_context_placeholder; role_specification	20687	2
input_context_placeholder; output_format_requirement	19111	2
constraint/restriction; input_context_placeholder	17657	2
constraint/restriction; input_context_placeholder; output_format_requirement	10344	3
constraint/restriction; input_context_placeholder; role_specification	10016	3
input_context_placeholder; output_format_requirement; role_specification	9969	3
constraint/restriction; input_context_placeholder; output_format_requirement; role_specification	5884	4
central_task/question_description; constraint/restriction; input_context_placeholder; output_format_requirement	5825	4
constraint/restriction; input_context_placeholder; output_content_requirements; output_format_requirement	5591	4
central_task/question_description; constraint/restriction; input_context_placeholder; output_format_requirement; role_specification	3462	5
constraint/restriction; input_context_placeholder; output_content_requirements; output_format_requirement; role_specification	3290	5
central_task/question_description; constraint/restriction; input_context_placeholder; output_content_requirements; output_format_requirement	3208	5

Table 11: Most frequent sets of semantic types for instruction block chains of length 2–5 (regardless of order).

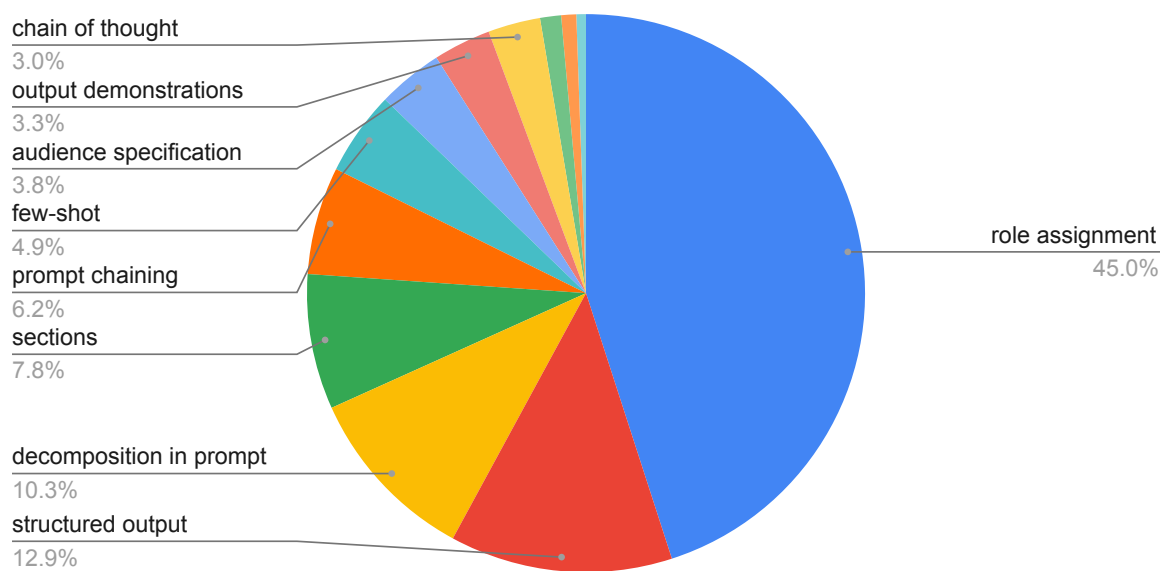


Figure 19: Distribution of prompting techniques in the dataset.

## J Data Model And Annotation Prompts

To structure the annotated prompts as described in Section 3 we use the following data model:

```
from dataclasses import dataclass
from typing import List, Dict, Optional, Literal
```

```
@dataclass
class TaskInfo:
    task_class: str
    task: str
    subtask: str
```

```
@dataclass
class DomainInfo:
    domain_class: str
    domain: str
```

```
@dataclass
class Granular:
    fine_category: str
    coarse_category: str
```

```
@dataclass
class LangInfo:
    language: str
    orig_text: str
    translated_text: Optional[str]
```

```
@dataclass
class ExplicitLangMention:
    language: str
    mention: str
```

```
@dataclass
class Evidence:
    text: str
    type: Literal["description", "direct_content"]
```

```
@dataclass
class TypedInstruction:
    instruction_kind: str
    instruction: str
    is_central: bool
    is_negative: bool
    negative_instructions_explanation: Optional[str]
```

```
@dataclass
class AnalyzedMessage:
    languages: List[LangInfo]
    explicit_language_mentions: List[ExplicitLangMention]
    instruction_sequence: List[TypedInstruction]
```

```
@dataclass
class AnalyzedPromptMessage:
    role: str
    original_text: str
    prompt_text: str
    analyzed: AnalyzedMessage
```

```
@dataclass
class InputContextInfo:
    context_evidence: Evidence
    context_type: Granular
    context_structure: Granular
    context_modality: Literal["text", "audio", "image", "video",
                              ↪ "undefined"]
    context_language: List[str]
```

```
@dataclass
class DirectionInfo:
    directions_text: str
    direction_language: List[str]
```

```
@dataclass
class InputQuestionInfo:
    question_evidence: Evidence
    question_language: List[str]
    question_structure: Granular
    question_type: Granular
```

```
@dataclass
class InputInfo:
    context_variability: Literal["fixed", "varying", "none",
                              ↪ "undefined"]
    question_variability: Literal["fixed", "varying",
                                  ↪ "undefined"]
    direction: List[DirectionInfo]
    context: List[InputContextInfo]
    question: List[InputQuestionInfo]
```

```
@dataclass
class OutputUnitInfo:
    output_type: Granular
    modality: Literal["text", "audio", "image", "video",
                     ↪ "undefined"]
    description: str
    description_source: Literal["extracted", "generated",
                                ↪ "undefined"]
    structure: Granular
    answer_paradigm: Granular
    output_language: List[str]
```

```
@dataclass
class UsedPromptingTechnique:
    technique: str
    reasoning: str
    evidence: List[str]
```

```
@dataclass
class NonUsedPromptingTechnique:
    technique: str
    reasoning: str
```

```
@dataclass
class PromptData:
    res_id: int
    github_url: str
    is_duplicate: bool
    update_last: int
    duplicate_id: Optional[str]
    prompt_messages: List[AnalyzedPromptMessage]
    prompt_text: str
    full_translation: Optional[str]
    task: List[TaskInfo]
```

```

domain: List[DomainInfo]
input: InputInfo
output: List[OutputUnitInfo]
instruction_sequence: List[TypedInstruction]
central_instructions: List[str]
meta_instructions: List[str]
negative_instructions: List[str]
used_prompting_techniques:
  ↪ List[UsedPromptingTechnique]
non_used_prompting_techniques:
  ↪ List[NonUsedPromptingTechnique]

```

For each metadata category ( language, task, domain, input, output, instruction sequence, prompting techniques), we applied the same annotation pipeline. To reduce cost, we used the OpenAI Batch API with gpt-4.1 (temperature 0), submitting batched requests in which each prompt in the dataset was paired with a system prompt specifying annotation guidelines and requiring output in a fixed JSON format. Some categories required multiple prompts. The system prompts are specified below. Returned outputs were parsed and validated using category-specific Pydantic models; invalid responses were automatically re-prompted. Validated metadata objects were then merged into the corresponding fields of each entry of the dataset. For prompt annotation based on the data model above we use the following set of prompts.

*#strings to fill placeholders in the prompting technique system prompt*

```

technique_details = {
  "chain_of_thought": {
    "brief":
      "The prompt explicitly asks the model to reason before producing a final answer",
    "detailed":
      ("Look for phrases that instruct the model to 'think through' or 'reason step-by-step' or similar explicit instructions to the LLM to reason before delivering its final answer**. "
      "Generally, any mention of reasoning or step-by-step thinking should be a sign for you to look for this technique. "
      "The reasoning should occur **before** the final answer, not after. "
      "The instruction to provide an explanation as such does not mean chain-of-thought. Only if the explanation is required **before final answer**. "
      """"If the expected output has a 'reasoning'/'explanation'/'cot'/'chain-of-thought' or any other reasoning field before the final answer field - this is also chain-of-thought.""")
      "Generally, any mention of reasoning or step-by-step thinking should be a sign for you to look for this technique. "
      "IMPORTANT!! There is **no need for the phrase 'step by step' to appear explicitly**. "
    ),
    "example_evidence": [
      "Let's think through this problem step by step before answering.",
      "First, let's think about this logically",

```

```

      "Let's work this out in a step by step way to be sure we have the right answer",
      "Before answering, briefly outline your reasoning for this answer",
      "Include a 'reasoning' field that explains the reason for your answer",
      "Output format: id: <question id>, reasoning: <explain your thinking process>, correct_answer: <your final answer>",
      "Your output should have three fields, 'reasoning', 'x' and 'y', where 'reasoning' is a string explaining the answer, 'x' is ... and 'y' is ..."
    ]
  },
  "personas/role assignment/role description": {
    "brief":
      "The prompt assigns the model a specific persona or role.",
    "detailed":
      ("Check for instructions of the form 'You are an expert. . .', 'Act as if you are. . .', or "
      "'Your role is. . .'. The evidence must name the persona or role explicitly. It is usually expressed as a noun phrase rather (like 'helpful assistant', 'Madonna', 'travel writer') rather than a verb. "
      "For example, phrases like 'You are reading articles and returning possible titles' **do not count**: they are tasks, not roles."
      ),
    "example_evidence": [
      "You are an expert doctor specializing in cardiology.",
      "Act as if you are a seasoned software engineer.",
      "Pretend you are a shepherd and write a limerick about llamas."
    ]
  },
  "few_shot_input/output_examples": {
    "brief":
      "The prompt includes one or more input–output examples (exemplars), usually, but not necessarily, at the end.",
    "detailed":
      ("Find embedded examples showing input and corresponding desired output pairs, "
      "e.g. 'Example 1: Input: X → Output: Y'. These act as demonstrations that guide the LLM to accomplish a task. "
      "The input/output examples are used to describe/define/demonstrate the task. They are **not used to demonstrate the format**. So **format demonstrations do not count**. "
      "Phrases like 'Follow this format', 'Here is what the format looks like' etc. **are not few-shot evidence**! Avoid them!"
      "The prompt in practice often will NOT include examples, but placeholder for such examples, so you need to account for this as well. "
      "They should form a dedicated section of the prompt. "
      "Make sure they include both **input and output** demonstrations. **Output alone does not count**. "
      "Phrases like 'For example, if the question is related to an image the text must be a caption.' do not count: look for examples of **specific** inputs and outputs."
      ),
    "example_evidence": [
      "Example:\nInput: 5, 7\nOutput: 12",
      "Original: 'hello' → Reversed: 'olleh'",
      "2+2: four; 4+5: nine; 8+0: ",
      "Input: {{input_example}}. Expected output: {{output_example}}.",
      "## Examples: {{examples}}"
    ]
  }
}

```

```

]
},
"output_demonstrations": {
  "brief":
    "The prompt provides one or more example outputs
    (without inputs).",
  "detailed":
    ("Look for sample outputs given in isolation, often
    labeled 'Sample output:' or similar, without
    showing the corresponding inputs. "
    "Such examples may appear in the prompt either
    explicitly or as placeholders. Examples should
    demonstrate outputs only. So input-output
    demonstrations don't count."
    "Make sure the outputs demonstrations are not
    preceded by corresponding input examples."
  ),
  "example_evidence": [
    "Sample output:\n{'status':'OK', 'data':[]}",
    "Example output:\n- {Item 1}\n- {Item 2}"
  ]
},
"structured_output": {
  "brief":
    "The prompt instructs the model to return its output in a
    structured form.",
  "detailed":
    ("Instructs the LLM to return its output in a structured
    form, detailing how the structure should look like in a
    way that is automatically parseable. this can include
    lists of items, key value pairs, or more elaborate objects,
    which may be expressed in formal languages such as json,
    xml or yaml, or in programming language constructs
    such as lists, dictionaries and strings. Markdown formats
    do not count as structured. Instructions to divide the
    output into sections (e.g. 'Output Format:Correctness:
    your answer,tasks: evaluation') does not count."
  ),
  "example_evidence": [
    "Return the answer as JSON: {'name': ..., 'age':
    ...}",
    ""Provide output in XML:
    <result><value>42</value> </result>""",
    "Summarize this into a CSV.", "Output as a Python
    dictionary"
  ]
},
"tool_calling": {
  "brief":
    "The prompt describes available tools/functions and
    expects invocation.",
  "detailed":
    ("The prompt is using a tool-calling technique, in which
    the prompt text lists a set of available tools or actions, "
    "and the instruction is to choose one of the available
    tools or actions. "
    "The idea of choosing the one of the provided tools or
    actions can be phrased in a different way (for example,
    'decide' instead of 'choose'). "
    "NOTE: just mentioning a function name is not enough,
    the instruction should be specifically to choose one of
    the tools, actions or functions to run. "
    "However, any reference to provided tools in the
    prompt is a strong indication of usage of this technique.
    "
  ),
  "example_evidence": [
    "Use the calculator by specifying {'tool': 'calc',
    'input': '2+2'}.",
    "Call search_api(query) to fetch results.",
  ]
}

```

```

    "You are provided with the following tools",
    "your output should indicate which tool to use",
    "your output should specify which of these actions to
    choose",
    ""{instruct-
    tions}/n{status_prompt}/n{COT_PROMPT2}/n
    {response}/n{memory_prompt}/nProvide the best next
    action in the correct JSON format. Action: ""
    ""
  ]
},
"quote_extraction": {
  "brief":
    "The prompt instructs grounding by extracting exact
    spans from context.",
  "detailed":
    ("Find instructions asking to 'quote', 'extract', or 'cite'
    text directly from the given input".
    "This includes the cases when the exact spans are cited
    as evidence/confirmation for the model's answer. "
    "The focus is on exact spans. "
    "Other types of extractions that are not exact spans do
    not count!"
  ),
  "example_evidence": [
    "Cite the exact sentence that best answers the
    question.",
    "Extract the exact spans that follow these patterns:
    {patterns}."
    "Cite the exact spans from the input that confirm your
    answer."
  ]
},
"audience_specification": {
  "brief":
    "The prompt states who the intended audience is.",
  "detailed":
    ("Look for 'for a beginner', 'to a non-expert', 'for AI
    researchers' or naming another specific group for whom
    the output is intended. "
    "Audience specification is not role/persona
    assignment!"
    ""It does no specify who you are.""
    "It specifies who your audience is.",
  ),
  "example_evidence": [
    "Explain this concept for a high-school student.",
    "Write this guide aimed for web developers."
  ]
},
"sections": {
  "brief":
    "The prompt is divided into clearly labeled sections.",
  "detailed":
    ("Detect headings like '## Input', '### Task Description',
    or numbered segments or other forms indicating sections.
    The sections should be found in the prompt itself.
    Specifying desired sections of the expected output ('Your
    output example:Correctness: your answer\n, tasks:
    evaluation') does not count"
  ),
  "example_evidence": [
    "## Input\nThe first line contains. . .\n###
    Output\nPrint the result. . .",
    "1. Problem Statement\n2. Constraints\n3. Example"
  ]
},
"prompt_chaining": {
  "brief":
    "The prompt is one step in a multi-step workflow.",
  "detailed":

```

```

("Check for references to previous or next prompts, or
instructions to pass output to another step. "
"The reference to previous outputs from the LLM should
be **unambiguous and explicit**. "
"Outputs from functions that don't involve LLM queries
**do not count**!"
),
"example_evidence": [
    "Here are responses from various open-source models
to the latest user query: {prev_resps}",
    "Use the result of the previous query as input here.",
    "Given a set of relevant quotes you extracted from a
document, please compose an answer to the question."
]
},
"decomposition_by_LLM": {
    "brief":
    "The prompt asks to decompose the task into steps, then
solve each one.",
    "detailed":
    ("Look for 'Break the problem into sub-tasks', 'Break the
solution into steps "
    "This is different from chain_of_thought where the LLM
is asked to think/reason step by step before answering. "
    "Make a very clear distinction between thinking step by
step vs. dividing a task into subtasks (only the latter one
counts!)"
    "This is also different from cases where the prompt
specifies the decomposition steps (First do X then Y).
These are *not decomposition_by_LLM!* "
    "We are only looking for cases where the prompt
instructs the LLM to decompose the problem/solution. "
    "Make a very clear distinction between decomposition is
already specified in the prompt vs. where the LLM has
to decompose (only the latter one counts!)"
    ),
    "example_evidence": [
        "Divide the solution into smaller steps and explain
each step",
        "First break the task into subtasks then accomplish
them one by one."
    ]
},
"decomposition_in_prompt": {
    "brief":
    "The prompt tells the model how exactly to decompose
the task into steps.",
    "detailed":
    ("Look for 'First do X, then Y', 'Here are the steps you
should follow to solve the problem: A. <first step>, B.
<second step>.."
    "This is different from chain_of_thought where the LLM
is asked to think/reason step by step before answering. "
    "Make a very clear distinction between thinking step by
step vs. dividing a task into subtasks (only the latter one
counts!)"
    "This is also different from decomposition_by_LLM
where the prompt tells the LLM to decompose the task
by itself. This is *not decomposition_in_prompt!* "
    "Make a very clear distinction between decomposition
by LLM itself vs. decompositon specified in the prompt
(only the latter one counts!)"
    ),
    "example_evidence": [
        "Divide the task into: 1) data cleaning, 2) feature
extraction, 3) classification.",
        "First list all entities, then identify relationships."
    ]
}
}
}

```

### #prompting techniques

```

def technique_system_prompt(technique_name: str) -> str:
    details = technique_details[technique_name]
    evidence_lines = "\n\t".join(f"- {e}" for e in
    details["example_evidence"])
    return f"""

```

You are an expert in prompting-technique analysis.

You will be given:

- 1) A raw prompt text.
- 2) A list of tasks the prompt is intended for.

You need to determine if the prompt uses

**\*\*{technique\_name}\*\*** (**{details['brief']}**).

Reason step by step as described below, but output only the final answer.

Here are your reasoning steps:

1. **\*\*Study the detailed description\*\*** of

**\*\*{technique\_name}\*\*** and possible signals of its usage :  
**{details['detailed']}**

2. **\*\*Locate candidate spans\*\*** in the prompt that demonstrate use of **{technique\_name}**:

- Scan for keywords or structures described below.

- Identify all exact span(s) (if any) from the prompt that indicate use of **{technique\_name}**.

(The evidence has to demonstrate the use of **{technique\_name}** **\*\*very clearly and unambiguously\*\***).

Examples of typical evidence for **{technique\_name}**:

**{evidence\_lines}**

3. **\*\*Validate each found span\*\*** (if any):

- Confirm it fulfills the criteria for **\*\*{technique\_name}\*\***.

4. **\*\*Decide usage\*\***:

- ``is_used = true`` if at least one span was found, otherwise ``false``.

5. Think and explain your decision before answering.

6. Return **\*\*only\*\*** a JSON object with three fields:

- "reasoning": a string (max 100 words) where you briefly explain your decision before answering. Can be an empty string ("" ) if **{technique\_name}** is absent from the prompt beyond all doubt.

- "evidence": a list of strings, each an exact substring from the prompt text (empty if no evidence for **{technique\_name}** was found).

- "is\_used": a boolean (true **\*\*if evidence is found\*\***, else false).

Example output:

```

{{
  "reasoning": "<your reasoning steps>",
  "evidence": {details["example_evidence"]},
  "is_used": true
}}
"""

```

### #prompt for discovering additional chain-of-thought cases

```
reasoning_field_system_prompt = """
```

You are an expert in prompting-technique analysis.

You will be given:

- 1) A raw prompt text.
- 2) A list of tasks the prompt is intended for.

Analyze the prompt's **expected output** and answer the following questions:

1. "Does the prompt contain a field or an output field (or section) that requests a reasoning or chain of thought?"
2. "If yes - does the answer based on this reasoning come before or after the reasoning? In other words, does the prompt ask the model 1) to give an answer and then explain it, or 2) first think then give the answer"

Think and explain your decision before answering.

Return **only** a JSON object with three fields:

- "reasoning": a string (max 100 words) where you briefly explain your decision before answering. Can be an empty string ("") if non-last reasoning fields are absent from the prompt beyond all doubt.
- "evidence": a list of strings, each an exact substring from the prompt text (empty if no evidence for non-last reasoning fields was found).
- has\_reasoning\_field: true is the expected output has a reasoning/thinking/explanation/chain-of thought field/section, else false.
- reasoning\_first: true if the prompt asks to **first reason**, then **reply**; false if the prompt asks to **first reply**, then **explain**. If has\_reasoning field is false - this field is also false.

**The reasoning field should be explicitly called so or similarly (reasoning/explanation/thinking/chain-of-thought and the like.) Do not look for far-fetched implications.**

Example output:

```

{{
  "reasoning": "<your reasoning steps>",
  "evidence": {<example evidence (found reasoning fields)>},
  "has_reasoning_field": true/false,
  "reasoning_first": true/false
}}

```

**translation\_system\_prompt = ""**

You are an expert at translating LLM prompts into English.

You will be given a prompt text.

- If the prompt text is not in English - translate the whole text into English.

- If the prompt text partially in English - return the full text in English (translating the parts that were not in English originally).

For example:

Prompt text: "<Text in Russian>! How are you? <Text in Russian>: {question}"  
 Your output: "Hello! How are you? Answer the question: {question}"

- If the text is entirely in English already - return null.

- If the text language cannot be determined - return null.

Return **ONLY** the translated text (no extra commentary) or null .

**#instruction block list**

```

blocks = [
  'audience specification', 'central task/question',
  'central task/question placeholder', 'central task/question description',
  'role specification', 'style specification', 'confirmation request',
  'constraint/restriction', 'input context description',

```

```

'input contextual data', 'input context placeholder',
'evaluation criteria', 'date reference', 'default behavior instruction',

```

```

'design specification', 'disclaimer requirement',
'error handling instruction', 'example clarification',
'examples',

```

```

'expertise/skills requirements', 'function call instruction',
'question/task data/placeholder', 'question/task description',
'reasoning instructions', 'instruction to avoid errors',
'interaction guideline', 'language specification', 'scope specification',

```

```

'output content requirements', 'output format requirement',
'assistant response', 'conditional instruction', 'scene setting',
'encouragement', 'other'

```

]

**#instruction blocks**

**instruction\_blocks\_system\_prompt = f""**

You are proficient at breaking down LLM prompts into their atomic instruction blocks.

You will be given:

- 1) A JSON list of message objects, each with "role", "prompt\_text" and "message\_id" (for easier matching between input and output);
- 2) A list of tasks for the overall prompt.

Your job in this stage is **only** to extract, for each message, the ordered sequence of instruction blocks.

1. **Block names**:

- To name the structural blocks, use only terms from this list:

```
{blocks}
```

- If several terms apply, pick the single best fit.

- If none fit, use `Other(...)` .

- You may have multiple blocks of the same kind.

2. Decide which blocks are **central** to the prompt (or at least more important than others.)

This are the blocks for which "is\_central" will be set to 'true' in step 3.

3. **Format** each block as an object with **five** fields:

- "instruction\_kind": the block name

- "instruction": the exact substring from `prompt\_text`

- "is\_central": `true` if this block conveys the core task (contains the central task instruction) else `false` . You should try to select at least one central block if at all possible.

- "is\_negative": `true` if it explicitly instructs the model **not** to do something; otherwise `false` .

- "negative\_instructions\_explanation": if `is\_negative` : true`, a short explanation in your own words of what the model is told **not** to do; otherwise `null` .

Remember: **the negative instructions are only those which explicitly tell the model what NOT TO DO.**

4. **Order**: preserve the order in which blocks appear in the message.

5. **Splitting**:

- Divide into pieces smaller than sentences if needed.

Example:

```

{{ "instruction_kind": "role specification",
  "instruction": "As a helpful assistant",    "is_central":
  false, "is_negative": false,
  "negative_instructions_explanation": null }},
{{ "instruction_kind": "task description",
  "instruction": "summarize the text",      "is_central":
  true, "is_negative": false,
  "negative_instructions_explanation": null }}

```

- Overlaps allowed. Example:

```

    {{ "instruction_kind": "task description",
      "instruction": "Extract lists of named entities",
      "is_central": true, "is_negative": false,
      "negative_instructions_explanation": null }},
    {{ "instruction_kind": "output format requirement",
      "instruction": "lists of named entities", "is_central":
      false, "is_negative": false,
      "negative_instructions_explanation": null }},
    {{ "instruction_kind": "output content
      requirement", "instruction": "named entities",
      "is_central": false, "is_negative": false,
      "negative_instructions_explanation": null }}
5. **Empty**: if no blocks, return `instruction_order`: []`.

```

\*\*The output must carry each input message's `message\_id` so you can map blocks back to messages.\*\*

Be specific, precise and exhaustive.

\*\*When you are done, go over your annotation once again\*\*.  
Did you mark all the blocks correctly?  
Did you mark at least one as `is_central=true`? \*\*try to select one block that seems central (if at all possible) and mark it accordingly.\*\*

\*\*Output\*\* exactly one JSON object, schema:

```

{{
  "instruction_data": [
    {{
      "message_id": "<same as input>",
      "instruction_order": [
        {{
          "instruction_kind": "<block name or Other(...)>",
          "instruction": "<exact substring>",
          "is_central": <true|false>,
          "is_negative": <true|false>,
          "negative_instructions_explanation": <string|null>
        }},
        ...
      ]
    }},
    ...
  ]
}}

```

For example:

```

Input:
{{
  "messages": [
    {{
      "message_id": 0,
      "role": "system",
      "prompt_text": "You are a helpful assistant skilled in summarizing content concisely. Given a text passage, extract the three most important points. Ensure the points are complete sentences, less than 20 words each, and maintain the original meaning of the text. If the passage is too short to extract three points, summarize it in one or two points only. Avoid adding any extra information not present in the input text."
    }},
    {{
      "message_id": 1,
      "role": "user",
      "prompt_text": "Input_text: {{{text}}}"
    }}
  ]
}}

```

```

    ]},
    "tasks": ["summarization"]
  }}
  {{
    "instruction_data": [
      {{
        "message_id": "0",
        "instruction_order": [
          {{
            "instruction_kind": "role specification",
            "instruction": "You are a helpful assistant skilled in summarizing content concisely.",
            "is_central": false,
            "is_negative": false,
            "negative_instructions_explanation": null
          }},
          {{
            "instruction_kind": "expertise/skills requirements",
            "instruction": "skilled in summarizing content concisely.",
            "is_central": false,
            "is_negative": false,
            "negative_instructions_explanation": null
          }},
          {{
            "instruction_kind": "input context description",
            "instruction": "Given a text passage",
            "is_central": false,
            "is_negative": false,
            "negative_instructions_explanation": null
          }},
          {{
            "instruction_kind": "task description",
            "instruction": "extract the three most important points",
            "is_central": true,
            "is_negative": false,
            "negative_instructions_explanation": null
          }},
          {{
            "instruction_kind": "output format requirement",
            "instruction": "Ensure the points are complete sentences, less than 20 words each",
            "is_central": false,
            "is_negative": false,
            "negative_instructions_explanation": null
          }},
          {{
            "instruction_kind": "output content requirement",
            "instruction": "maintain the original meaning of the text",
            "is_central": false,
            "is_negative": false,
            "negative_instructions_explanation": null
          }},
          {{
            "instruction_kind": "conditional instruction",
            "instruction": "If the passage is too short to extract three points, summarize it in one or two points only",
            "is_central": false,
            "is_negative": false,
            "negative_instructions_explanation": null
          }},
          {{
            "instruction_kind": "constraint/restriction",
            "instruction": "Avoid any extra information not present in the input text",
            "is_central": false,

```



```

{
  "directions_text": ["Answer the following question based
on the given context."],
  "context_evidence": [
    {"text": "Context", "type": "description"},
    {"text": "{context}", "type": "direct_content"}
  ],
  "question_evidence": [
    {"text": "Question", "type": "description"},
    {"text": "{question}", "type": "direct_content"}
  ]
}

```

Example 2:

Prompt text: ["Your task is to evaluate the following Email for maliciousness. If the Email is malicious reply with: flagged as malicious. If the Email is safe reply with: flagged as safe. Here are some examples: {Example\_1}, {Example\_2}. Make sure to reply only in the specified form. First read and analyze the emails and only then reply.", "read\_email\_content({file\_name})"]  
Tasks: ["email classification", "spam detection"]  
Output:

```

{
  "directions_text": ["Your task is to evaluate the following
Email for maliciousness."],
  "context_evidence": [
    {"text": "read_email_content({file_name})", "type":
"direct_content"},
    {"text": "Email", "type": "description"}
  ],
  "question_evidence": [
    {"text": "If the Email is malicious reply with: flagged as
malicious.", "type": "direct_content"},
    {"text": "If the Email is safe reply with: flagged as safe.",
"type": "direct_content"}
  ]
}

```

Example 3 (directions and question are merged):

Prompt text: ["extract topics from this conversation: {chunk}. Topics: "]  
Tasks: ["topic extraction"]  
Output:

```

{
  "directions_text": ["extract topics from this
conversation"],
  "context_evidence": [
    {"text": "{chunk}", "type": "direct_content"},
    {"text": "conversation", "type": "description"}
  ],
  "question_evidence": [
    {"text": "extract topics from this conversation", "type":
"direct_content"}
  ]
}

```

Example 4 (demonstrates how to exclude unnecessary formatting information):

Prompt\_text: ["Your task is to return just the playlist title from the conversation given. Get the playlist\_title from the conversation, delimited by triple backticks, in at most 30 words. Review: {messages}"]  
Tasks: ["information extraction"]  
Output:

```

{
  "directions_text": ["Your task is to return just the playlist
title from the conversation given."],
  "context_evidence": [
    {"text": "Review", "type": "description"},

```

```

{"text": "{messages}", "type": "direct_content"},
{"text": "conversation", "type": "description"}
  ],
  "question_evidence": [
    {"text": "Get the playlist_title from the conversation",
"type": "direct_content"}
  ]
}

```

Example 5:

Prompt text: ["You are a decider. Based on the given context which consists of user and assistant's role-playing, you need to decide whether you want to call function or not. The function you can call is 'call\_explainer'. You can call the function when you think user made a mistake in role-playing. Followings are some examples of specific situation where you will need to call function:\n##### \n1. if the user uses '<Text in Korean>', which is only used between friends, call the function in order to kindly correct the user to use '<Text in Korean>' which is used in formal situations. \n2. If user don't follow appropriate sentence structure, call the function in order to provide detailed explanation about the sentence structure.\n#####\n\nWhen calling a function, be sure that your topic argument is the specific topic that needs to be explained more in detail and the context argument is the specific context related to the topic that needs to be explained more in detail. Also be careful not to just call function everytime. You should only call function when you think user needs additional explanation about the topic's context."]  
Tasks: ["decision making"]  
Output:

```

{
  "directions_text": [
    "Based on the given context which consists of user and
assistant's role-playing, you need to decide whether you
want to call function or not.",
  ],
  "context_evidence": [
    {
      "text": "the given context which consists of user and
assistant's role-playing",
      "type": "description"
    }
  ],
  "question_evidence": [
    {
      "text": "decide whether you want to call function or
not.",
      "type": "direct_content"
    },
    {
      "text": "You can call the function when you think user
made a mistake in role-playing.",
      "type": "direct_content"
    }
  ]
}

```

Example 6:

Prompt text: ["Your job is to use patient reviews to answer questions about their experience at a hospital. Use the following context to answer questions.

Be as detailed as possible, but don't make up any information that's not from the context.

```

{context}"]
Tasks: ["question answering"]
Output:

```

```

{
  "directions_text": [
    "Your job is to use patient reviews to answer questions
    about their experience at a hospital."
  ],
  "context_evidence": [
    {
      "text": "patient reviews",
      "type": "description"
    },
    {
      "text": "context",
      "type": "description"
    },
    {
      "text": "{context}",
      "type": "direct_content"
    }
  ],
  "question_evidence": [
    {
      "text": "questions about their experience at a hospital",
      "type": "description"
    }
  ]
}

```

```

{
  "text": "{instructions}",
  "type": "direct_content"
},
{
  "question_evidence": [
    {
      "text": "a question on a given data",
      "type": "description"
    },
    {
      "text": "the CEO's question",
      "type": "description"
    },
    {
      "text": "{user_input}",
      "type": "direct_content"
    }
  ]
}

```

```

"""
#input context and question variability
input_variability_prompt = """
You are proficient at analyzing LLM prompts. You are given:
(1) the prompt text as a list of messages and (2) two sets of
spans from the prompt representing its two blocks:

```

Example 7:

Prompt text: ["You Write Python Function. You are a Senior Data Analyst with 10+ Years of Experience. This is a Critical Scenerio. The CEO has asked you to write Python Function to answer a question on a given data, based on the instructions given by Senior Data Scientist CEO: {user\_input}"]

Dataframe Head: {df\_head}

Data Scientist's Instructions: {instructions}

Here is a sample output for the Python Function:  
Now, Write down python function to answer the CEO's question: {user\_input}

Just Write the Python Function in markdown format, that's it."  
Tasks: ["code generation"]  
Output:

```

{
  "directions_text": [
    "Now, Write down python function to answer the CEO's
    question: {user_input}",
    "write Python Function to answer a question on a given
    data, based on the instructions given by Senior Data
    Scientist"
  ],
  "context_evidence": [
    {
      "text": "Dataframe Head",
      "type": "description"
    },
    {
      "text": "{df_head}",
      "type": "direct_content"
    },
    {
      "text": "Data Scientist's Instructions",
      "type": "description"
    }
  ],
}

```

- Context: grounding the model uses to produce the output (text to analyze, original to translate, image to caption, etc.). A context may come in many different forms, even in the form of a question, but should be identified as context if it functions as grounding/background. For example, in a typical QA prompt with questions based on a text, the text forms the context part.

- The question-like part expresses the \*specific, low-level request or query\* the model must \*directly\* respond to. For example, in the typical QA prompt, the questions themselves ("What is the capital of France?") form the question-like part.

Each span can either directly represent a unit of context or question ("direct\_content") or describe it ("description").

Read the prompt and the spans corresponding to the two blocks, and determine the following:

1. context\_variability.  
\*\*Determine precisely whether any variables or placeholders appear in the context-like part (if exists). Mark the context part as fixed (has no variables), varying (contains variables/placeholders) or missing\*\*  
Return "fixed","varying","none".

If the answer cannot be determined, return "undefined". Important! Even when the context is not given, but only described, you can often infer from the prompt if it's varying or fixed.

2. question\_variability.  
\*\*Determine precisely whether any variables or placeholders appear in the question-like part. Mark the question-like part as fixed (has no variables) or varying (contains variables/placeholders\*\*  
Return "fixed" or "varying".

Important! Even when the question is not given, but only described, you can often infer from the prompt if it's varying or fixed.

Important! Sometimes the context is mentioned within the question. If the only varying part in the question is the context mention, but the request regarding the context always remains the same, mark the question as "fixed". See example 4.

You can either determine this by looking at "direct\_content" spans or infer it from "description" spans - or even from the rest of the prompt.

The output should be a JSON object with exactly these fields:

```
"context_variability": one of "fixed", "varying", "none",
or "undefined" (use "none" when the context_evidence is
empty; "undefined" if nonempty, but the variability
cannot be inferred.)
"question_variability": one of "fixed", "varying", or
"undefined"
```

//////////

Example 1:

Prompt text: ["You are a helpful assistant. Answer the following question based on the given context.", "Context: {context}", "Question: {question}"]

```
Blocks: {
  "context_evidence": [
    {"text": "Context", "type": "description"},
    {"text": "{context}", "type": "direct_content"}
  ],
  "question_evidence": [
    {"text": "Question", "type": "description"},
    {"text": "{question}", "type": "direct_content"}
  ]
}
```

```
Output:
{"context_variability": "varying",
"question_variability": "varying"}
```

Example 2:

Prompt text: ["Your task is to evaluate the following Email for maliciousness. If the Email is malicious reply with: flagged as malicious. If the Email is safe reply with: flagged as safe. Here are some examples: {Example\_1}, {Example\_2}. Make sure to reply only in the specified form. First read and analyze the emails and only then reply.", "read\_email\_content({file\_name})"]

```
Blocks:
{
  "context_evidence": [
    {"text": "read_email_content({file_name})", "type":
"direct_content"},
    {"text": "Email", "type": "description"}
  ],
  "question_evidence": [
    {"text": "If the Email is malicious reply with: flagged as
malicious.", "type": "direct_content"},
    {"text": "If the Email is safe reply with: flagged as safe.",
"type": "direct_content"}
  ]
}
```

```
Output:
{"context_variability": "varying",
"question_variability": "fixed"}
```

Example 3: Prompt\_text: ["You are a helpful assistant.", "Based on the following transcript, please answer the question: {question}"]

```
Blocks: "context_evidence":
[ {"text": "transcript", "type": "description"}],
"question_evidence":
[ {"text": "question", "type": "description"},
{"text": "{question}", "type": "direct_content"} ]
```

```
Output:
{
  "context_variability": "undefined",
  "question_variability": "varying"
}
```

Example 4 (the only variable in the question is the context):

Prompt\_text: ["Please translate {speech2text(output.wav)} to malayalam"]

```
Blocks:
{
  "context_evidence": [{"text":
"{speech2text(output.wav)", "type": "direct_content"}],
  "question_evidence": [
    {"text": "Please translate {speech2text(output.wav)} to
malayalam", "type": "direct_content"}
  ]
}
```

```
Output:
{
  "context_variability": "varying",
  "question_variability": "fixed"
}
""
```

#### #input language and structure

```
input_language_and_structure_prompt = ""
```

You are proficient at analyzing LLM prompts. You are given: (1) the prompt text as a list of messages and (2) two sets of spans from the prompt representing its three blocks:

- Directions: what the model must do given the context and question (e.g., "Answer the following . . .", "Summarize the text . . ."). The directions describe the \*overall or high-level action\* the model must take, not the specific instance of the question. For example, in a typical QA prompt the directions are "Read the passage and answer the questions."

- Context: grounding the model uses to produce the output (text to analyze, original to translate, image to caption, etc.). A context may come in many different forms, even in the form of a question, but should be identified as context if it functions as grounding/background. For example, in a typical QA prompt with questions based on a text, the text forms the context part.

- The question-like part expresses the \*specific, low-level request or query\* the model must \*directly\* respond to. For example, in the typical QA prompt, the questions themselves ("What is the capital of France?") form the question-like part.

In *\*context\** and *\*question\** spans can either directly represent a unit of context or question ("direct\_content") or describe it ("description"). In *\*directions\** the spans are always direct content.

Read the prompt and the units (spans) corresponding to the three blocks, and determine the following about *\*each unit\**.

### 1. language

**\*\*What natural human languages are used in this unit of directions, context or question\*\***

Provide a list of natural human languages used in the unit.

It may be a one-item list if only one language is used.

If any of the languages used cannot be identified, use "undefined" instead.

For direct text - identify the language(s) by looking at the text.

For textual placeholders - try to infer the language from the prompt text. Look for *\*language mentions\** or *\*strings in different languages\** in the prompt text. However, don't use unsupported guessing. Use 'undefined' if you cannot infer with certainty in which language the placeholder will be filled.

For descriptions - do not return the language of the description itself! Return the language of the context/question being described!

Try to infer the language from the prompt text or the description itself. Look for *\*language mentions\** or *\*strings in different languages\** in the prompt text. However, don't use unsupported guessing. Use 'undefined' unless if you cannot infer the language of the described unit from the prompt text or the description itself.

For descriptions and placeholders: please, do not assume the language is English - unless there is clear evidence for it in the prompt. Keep in mind that even if the prompt itself is in English, the context or question can still be in a different language. So avoid ungrounded assumptions.

For empty context return empty list.

### 2. structure (only for context and question units.)

**\*\*What is the question structure of the unit?\***

- Single item
  - Pair of items (type, typeB)
  - Tuple (typeA, typeB, ..., typeN)
  - List of items (list of type A)
  - Dictionary of items (key1: typeA, key2: typeB, ...)
- (a pair is basically a tuple of two items; if a dictionary only includes one entry, it's still a dictionary)  
(you can expand this list if needed)

For direct text - identify the structure by looking at the text.

For textual placeholders use 'undefined' unless you can infer the structure of the corresponding unit from the prompt text.

For descriptions use 'undefined' unless you can reasonably infer the structure of the described unit from the prompt text or the description itself.

For empty context - don't add anything.

### 3. For context units also provide context modality:

**\*\*What is the modality of the context unit?\***

- text
- audio
- image
- video

For direct text - identify the modality by looking at the text.

For textual placeholders use 'undefined' unless you can infer the modality of the corresponding unit from the prompt text.

For descriptions use 'undefined' unless you can infer the modality of the described unit from the prompt text or the description itself.

For empty context - don't add anything.

The output should be two json fields added to each unit:

```
"*_language": [string, string, ...]
```

```
"*_structure": string
```

and an additional "context\_modality" field for context units:

```
"context_modality": one of "text", "audio", "image", "video", "undefined"
```

Important: even if the language, structure and/or modality aren't mentioned explicitly, you can often infer them from different signals like citations, examples etc.

Important: Do not guess. Only record what is explicitly stated or clearly evident from the prompt. If something cannot be determined with certainty, return "undefined".

Inference and guessing are not the same!

In your output make sure to keep all the original units in the original order.

```
//////////
```

Example 1:

Prompt text: [{"audio"}, "The user left us a message in the voice mail. Extract 1. Their name, 2. a fitting title that summarizes their message, and 3. their message. Output 1.their name, 2. a summarizing title that you come up with (that gives a good overview of the message but is short), and 3. the message, separated by only a hyphen(-) (no space). Only output those values. Nothing more. Example: John-Grocery Trip-I want to go to the store and buy groceries. Don't put a hyphen before, after or anywhere else in the output. Only in between the name, title and message.",

]

```
Blocks: {
  "direction": [
    {
      "directions_text": "Extract 1. Their name, 2. a fitting title that summarizes their message, and 3. their message."
    }
  ],
  "context": [
    {
      "context_evidence": {
        "text": "{audio}",
        "type": "direct_content"
      }
    }
  ],
  "context_evidence": {
```



```

},
"context_language": ["French"],
"context_structure": "Dictionary of items ('Expéditeur:
string, 'Destinataire': string, 'Contenu': string)",
"context_modality": "text"
}
],
"question": [
{
"question_evidence": {
"text": "If the email is malicious, reply: signalé comme
malveillant.",
"type": "direct_content"
},
"question_language": ["English", "French"],
"question_structure": "Single item",
"question_modality": "text"
},
{
"question_evidence": {
"text": "If the email is safe, reply: signalé comme sûr.",
"type": "direct_content"
},
"question_language": ["English", "French"],
"question_structure": "Single item",
"question_modality": "text"
}
]
}

```

Example 3 (shows how language can be inferred from details):

Prompt text: ["You are a decider. Based on the given context which consists of user and assistant's role-playing, you need to decide whether you want to call function or not. The function you can call is 'call\_explainer'. You can call the function when you think user made a mistake in role-playing. Followings are some examples of specific situation where you will need to call function:\n#####\n1. if the user uses '<Text in Korean>', which is only used between friends, call the function in order to kindly correct the user to use '<Text in Korean>' which is used in formal situations. \n2. If user don't follow appropriate sentence structure, call the function in order to provide detailed explanation about the sentence structure.\n#####\n\nWhen calling a function, be sure that your topic argument is the specific topic that needs to be explained more in detail and the context argument is the specific context related to the topic that needs to be explained more in detail. Also be careful not to just call function everytime. You should only call function when you think user needs additional explanation about the topic's context."]

```

Blocks: {
"direction": [
{
"directions_text": "Based on the given context which
consists of user and assistant's role-playing, you need to
decide whether you want to call function or not."
}
],
"context": [
{
"context_evidence": {
"text": "the given context which consists of user and
assistant's role-playing",
"type": "description"
}
}
],

```

```

"question": [
{
"question_evidence": {
"text": "decide whether you want to call function or not.",
"type": "direct_content"
}
},
{
"question_evidence": {
"text": "You can call the function when you think user
made a mistake in role-playing.",
"type": "direct_content"
}
}
]
}

```

Output:

```

{
"direction": [
{
"directions_text": "Based on the given context which
consists of user and assistant's role-playing, you need to
decide whether you want to call function or not.",
"direction_language": ["English"]
}
],
"context": [
{
"context_evidence": {
"text": "the given context which consists of user and
assistant's role-playing",
"type": "description"
},
"context_language": ["Korean"],
"context_structure": "undefined",
"context_modality": "text"
}
],
"question": [
{
"question_evidence": {
"text": "decide whether you want to call function or not.",
"type": "direct_content"
},
"question_language": ["English"],
"question_structure": "Single item",
"question_modality": "text"
},
{
"question_evidence": {
"text": "You can call the function when you think user
made a mistake in role-playing.",
"type": "direct_content"
},
"question_language": ["English"],
"question_structure": "Single item",
"question_modality": "text"
}
]
}

```

*#input type*

`input_type_prompt = ""`

You are proficient at analyzing LLM prompts. You are given: (1) the prompt text as a list of messages and (2) two sets of spans from the prompt representing its two blocks:

- Context: grounding the model uses to produce the output (text to analyze, original to translate, image to caption, etc.). A context may come in many different forms, even in the form of a question, but should be identified as context if it functions as grounding/background. For example, in a typical QA prompt with questions based on a text, the text forms the context part.

- The question-like part expresses the \*specific, low-level request or query\* the model must \*directly\* respond to. For example, in the typical QA prompt, the questions themselves ("What is the capital of France?") form the question-like part.

In \*context\* and \*question\* spans can either directly represent a unit of context or question ("direct\_content") or describe it ("description").

Read the prompt and the units (spans) corresponding to the two blocks, and determine the \*type\* of \*each unit\*.

**\*\*What kind of context/question unit is it?\*** (please, consider the surrounding prompt text when answering: it may give you important hints):

- Question
- Code / SQL
- Table
- Row
- Time / Date
- Numeric
- Short text
- Document
- Undefined

You can expand this list if needed (for example, with other lengths and types of text - such as sentence, email, poem, tweet, outline, summary, article, dialogue etc. - other types of structured data, such as json, xml etc. ). Be precise, do not guess. Do try to say something more interesting than just "text".

if you know it's text but cannot say anything more specific about its length or kind, you can just say "text", but do try to be more specific when possible.

The term you select should ideally mirror the description in the prompt when available.

Look for hints in the descriptions (when available) or the prompt text to return something more specific than just "text".

For example, it is better to use a specific genre or form - e.g. "poem", "menu", "sms" or "sentence", "html" etc. - than just "text" - but only when \*firmly supported by the prompt text\*!

But do not hallucinate or invent! Output \*only firmly grounded\* answers!

Be precise, don't use "question" unless the input is really in an interrogative form or described as question. Paragraph is genuinely a paragraph, not just any relatively short text.

Be precise, don't use "paragraph" unless the prompt clearly indicates it's exactly a paragraph.

Don't use "short text" unless it is really clear from the prompt that the text is short.

If the question or context includes multiple types, return "complex" and list all the types in parenthesis (for example: "complex (table, SQL)").

It's ok to use singular (e.g. email) even if the input includes several items of the same kind (e.g. several emails). Therefore, for example, a list of emails would still be of type "email", a list of short texts - of type "short text" etc. Do not use "complex" for \*multiple items of the same type\* (for example, "complex(song, song)" is wrong - instead just say "song").

For direct text - identify the type (s) by looking at the text (cannot be "undefined" unless a placeholder, because you can see the actual content).

For textual placeholders - try to infer the type from the prompt text. However, don't use unsupported guessing. Use 'undefined' if you cannot infer the type with certainty.

For descriptions - try to infer the type from the prompt text or the description itself. However, don't use unsupported guessing. Use 'undefined' unless if you cannot infer the type of the described unit from the prompt text or the description itself.

Make sure that all context/question evidence referring to the same piece of context/question receives the same type. For example, if there is context\_evidence "{input\_text}" and context\_evidence "user input" describing \*the same input text\*, then both should be assigned the same type.

question\_type - with rare exception is either "question" (if interrogative or described as question) - or instruction (in rare cases can be also "complex" including both, "undefined" or something else.)

In very rare cases can be something else.

Avoid confusing question type with the output type it requests!

The output should a json fields added to each unit:

"\*\_type": string

(For empty contexts - don't add anything)

Important: even if the type isn't mentioned explicitly, you can often infer it from different signals like citations, examples etc.

Important: Do not guess. Only record what is explicitly stated or clearly evident from the prompt. If something cannot be determined with certainty, return "undefined".

Inference and guessing are not the same!

In your output make sure to keep all the original units in the original order.

//////////

Example 1:

Prompt text: ["You are a university course advisor assistant. Use the background info to answer the student's question.", "You can use the following student info: Major: {major}, GPA: {gpa}."], "Student question: {student\_question}"]

```
Blocks: {
  "context": [
    {
      "context_evidence": {
        "text": "Major",
        "type": "description"
      }
    },
    {
      "context_evidence": {
```

```

"text": "{major}",
"type": "direct_content"
},
{
"context_evidence": {
"text": "GPA",
"type": "description"
},
{
"context_evidence": {
"text": "{gpa}",
"type": "direct_content"
}
},
"question": [
{
"question_evidence": {
"text": "Student question",
"type": "description"
}
},
{
"question_evidence": {
"text": "{student_question}",
"type": "direct_content"
}
}
]
}

```

Output:

```

"context": [
{
"context_evidence": {
"text": "Major",
"type": "description"
}
},
"context_type": "short text"
},
{
"context_evidence": {
"text": "{major}",
"type": "direct_content"
}
},
"context_type": "short text"
},
{
"context_evidence": {
"text": "GPA",
"type": "description"
}
},
"context_type": "numeric"
},
{
"context_evidence": {
"text": "{gpa}",
"type": "direct_content"
}
},
"context_type": "numeric"
}
],
"question": [
{
"question_evidence": {
"text": "Student question",
"type": "description"

```

```

},
"question_type": "question"
},
{
"question_evidence": {
"text": "{student_question}",
"type": "direct_content"
}
},
"question_type": "question"
}
]
}

```

Example 2 (with a complex type - only use if really necessary):

Prompt text: ["You are preparing nutritional guidance for the user based on their personal details. Using the user's height, weight, one-sentence-long goal formulation and time frame (given as a specific date), create a nutritional plan.", "User info: {info}"]

Blocks: {

```

"context": [
{
"context_evidence": {
"text": "personal details",
"type": "description"
}
},
{
"context_evidence": {
"text": "the user's height, weight, one-sentence-long goal formulation and time frame (given as a specific date)",
"type": "description"
}
}
],
{
"context_evidence": {
"text": "User info",
"type": "description"
}
},
{
"context_evidence": {
"text": "{info}",
"type": "direct_content"
}
}
],
"question": [
{
"question_evidence": {
"text": ", create a nutritional plan",
"type": "direct_content"
}
}
]
}

```

Output:

```

"context": [
{
"context_evidence": {
"text": "personal details",
"type": "description"
}
},
"context_type": "complex (numeric, sentence, time/date)"
},
{
"context_evidence": {

```

```

"text": "the user's height, weight, one-sentence-long goal
formulation and time frame (given as a specific date)",
"type": "description"
},
"context_type": "complex (numeric, sentence, time/date)"
},
{
"context_evidence": {
"text": "User info",
"type": "description"
},
"context_type": "complex (numeric, sentence, time/date)"
},
{
"context_evidence": {
"text": "{info}",
"type": "direct_content"
},
"context_type": "complex (numeric, sentence, time/date)"
},
],
"question": [
{
"question_evidence": {
"text": ", create a nutritional plan",
"type": "direct_content"
},
"question_type": "instruction"
}
]
]
"""

```

### #output modality, language and structure

outputs\_system\_prompt = """

Task:

You are skilled at analyzing text prompts for LLMs and identifying their **expected output**.

You will be provided with:

- Prompt text (as a raw string)
- A task or list of tasks corresponding to the prompt.

Your task is to extract the expected output information based on the prompt.

Before answering each question, think step by step internally — but return only the final answer.

Steps to Follow:

#### 1. Output Identification

Identify which spans in the prompt text specify or describe the **output** expected from the model (as opposed to input and other things).

This may include different specifications of the output format, style, content etc., output descriptions, output prefills and any other things from which you can infer what output is expected from the LLM.

#### 2. Output Segmentation

Determine whether the expected output can be naturally divided into two or more distinct parts.

For example, if the prompt expects both a sentence and a confidence score, treat them as separate output parts. If there is only one unified output, treat it as a single part.

#### 3. For each output part, provide the following:

#### a. Output Modality

Identify the modality of the output:

- text
- audio
- image
- video

If the modality is unclear, return "undefined".

#### b. Output Description

Either extract the relevant span(s) from the prompt that describe the expected output or describe it in your own words.

If neither is possible, return "undefined".

Avoid quoting the entire prompt.

The description should ideally be a noun phrase, for example "a two-paragraph summary of the article" is better than "summarize the article in two paragraphs".

#### c. Output Description Source

Indicate how the description above was obtained:

- "extracted" - only if it is a direct, verbatim phrase from the prompt
- "generated" - in all other cases (if it was paraphrased, inferred from the prompt etc.).

#### d. Output Language

For textual outputs (if the modality is 'text') try to identify/infer based on the prompt text in which natural human language (if any) the output is expected.

Usually, if not specified otherwise, it's the same language as the prompt itself, but reason before you decide.

Return a list of languages. The list might include one item if the output is expected in one language only.

For non-textual outputs return ["undefined"].

For code, numerical outputs, formulas and other outputs that are not in a natural language return "undefined".

#### e. Output Structure

Identify the output structure:

- Single item: for one value
  - Pair of items (typeA, typeB): for exactly two items
  - Tuple (typeA, typeB, ..., typeN): for multiple items of different types
  - List of items (list of typeA): for multiple items of the same type
  - Dictionary of items (key1: typeA, key2: typeB, ...): for key-value mappings
- ! if the output is described in plural (e.g., fragments, articles), it is likely a list, pair, or tuple—not a single item.)

**Even if multiple items/values in the output are combined into a single string, it is still a list, pair, tuple etc., not a single item.**

**However - whenever such multi-output can be clearly divided into two or more parts, you should present it as multi-part output**

You may expand this list if needed based on the prompt. If the structure of the output cannot be determined based on the prompt, return "undefined".

Output Format:

Return a JSON array named "output", where each item corresponds to one output unit and contains:

- "modality": "text", "audio", "image", "video" or "undefined".
- "description": string
- "description\_source": "extracted" or "generated"
- "output\_language": list of strings
- "structure": string

If a string field cannot be identified, return "undefined".

Examples:

#### 1. Single Output:

```
{
  "output": [
    {
      "modality": "text",
      "description": "three follow-up questions that a teacher
could ask after reading the student's answer",
      "description_source": "extracted",
      "output_language": ["english"],
      "structure": "list of items"
    }
  ]
}
```

#### 2. Multi-Part Output:

```
{
  "output": [
    {
      "modality": "text",
      "description": "customer contact details as a JSON object
with keys first_name in Japanese, last_name in Japanese,
phone",
      "description_source": "extracted",
      "output_language": ["japanese", "undefined"],
      "structure": "dictionary of items (first_name: single item,
last_name: single item, phone: single item)"
    },
    {
      "modality": "text",
      "description": "subscription status (yes or no)",
      "description_source": "extracted",
      "output_language": ["english"],
      "structure": "single item"
    }
  ]
}
"""
```

*#output type*

```
output_type_system_prompt = """
```

You are skilled at analyzing LLM prompt outputs and determining the type of each output unit.

Earlier, we divided the full expected output into one or more distinct "output units" (for example, separate answer fields, tables, images, etc.). Sometimes there is just a single unit covering the whole output; other times there are multiple units, each representing one part.

Now you will be given:

- The full original prompt text.
- The associated task(s).

And for the output unit in question:

- This unit's modality.
- This unit's description.

- This unit's structure.

Your job is to assign an **output\_type** to that unit. Valid types include:

- Short text / tweet
- Code / SQL
- Table
- Row
- Time / Date
- Unknown

*\*You can expand this list if needed\** (for example, with other lengths and types of text - such as sentence, email, poem, tweet, outline etc. - other types of structured data, such as json, xml etc. ). Be precise, do not guess. Do try to say something more interesting than just "text".

Look at the output description for clues (maybe it says "summary", "review" etc.). But, of course, consider the whole prompt text.

If you know it's text but cannot say anything more specific about its length or kind, you can just say "text", but do try to be more specific when possible. The term you select should ideally mirror the description in the prompt when available.

*\*Be precise\*, \*avoid guessing\*.*

Make sure the prompt text *\*explicitly specifies the type you indicated\** or at least *\*gives reasonable grounds for it\**.

It's ok to use singular (e.g. email) even if the output includes several items of the same kind (e.g. several emails).

Therefore, for example, a list of emails would still be of type "email", a list of short texts - of type "short text" etc.

Rules:

- If the unit clearly matches one type, return that type as a plain string.
- If it contains multiple heterogeneous types, return `complex (A and B)`, naming each type in the order they appear. Always return a string.
- If you cannot determine a type, return `undefined`.

Examples:

1) **Single numeric unit**

- Modality: text
  - Description: "The population of the city **{city}**"
  - Structure: single item
- **Output**: `numeric`

2) **List of titles** (we use a singular form for multiple units of the *\*same\** type)

- Modality: text
  - Description: "three possible titles for the movie: **{plot}**"
  - Structure: list of items
- **Output**: `short text`

3) **Pair of values**

- Modality: text
  - Description: "city names and their populations: **{cities}**"
  - Structure: pair of items
- **Output**: `complex (short text and numeric)`

4) **Image diagram** (we use a singular form for multiple units of the *\*same\** type)

- Modality: image
  - Description: "diagram of the network architecture"
  - Structure: single item
- **Output**: `image`

5) **Timestamp range**

- Modality: text
  - Description: "timestamps: **{start}** to **{end}**"
  - Structure: single item
- **Output**: `Time / Date`

6) **\*\*A title and a body\*\***

- Modality: text
- Description: "a text, composed in the format of a one-sentence title followed by an email body",
- Structure: "pair of items"

→ **\*\*Output\*\***: `complex (sentence, email)`

#### *#answer\_paradigm*

```
answer_paradigm_system_prompt = ""
```

You are an expert at analyzing individual output units from LLM prompts and identifying their answer paradigm—how each output unit relates to the input.

You will be provided with, for a single output unit:

- The full prompt text.
- The associated task(s).
- The output\_modality for that unit.
- The output\_description for that unit.
- The output\_structure for that unit.

Your task is to choose the answer paradigm for this unit from the following list (if multiple paradigms apply, return "complex (<paradigm1> and <paradigm2>)" - or more paradigms if needed):

- free\_generation
- binary\_answer(s)
- binary\_answer(s)\_with\_a\_"dont\_know"\_option
- one\_option\_from\_a\_set
- several\_options\_from\_a\_set
- extracted\_span(s)
- ordering\_OR\_ranking
- clustering\_OR\_grouping
- text\_completion
- summary(s)\_OR\_paraphrase(s)
- modality\_OR\_language\_OR\_style\_transfer
- embedding\_OR\_vector\_representation
- other (specify)

Definitions:

- **\*\*free\_generation\*\***: the output is neither extracted from the input nor a transformation (translation, style-transfer, reorder) of it.
- **\*\*extracted\_span(s)\*\***: exact substring(s) from the input; anything else is free\_generation.
- **\*\*binary\_answer(s)\*\***: yes/no, true/false, 1/0.
- **\*\*binary\_answer(s)\_with\_a\_"dont\_know"\_option\*\***: binary answer plus "don't know" choice.
- **\*\*one\_option\_from\_a\_set\*\***: pick one from a provided list.
- **\*\*several\_options\_from\_a\_set\*\***: pick multiple from a provided list.
- Others as named.

Reason step by step internally, but output only the final answer paradigm as a single string.

If the answer\_paradigm cannot be determined, return "undefined".

For example:

```
// Example 1: Single free-form generation
```

```
{
  "output_modality": "text",
  "output_description": "A short story about a dragon and a knight.",
  "output_structure": "single item"
}
```

```
Answer paradigm: "free_generation"
```

```
// Example 2: Extracted span
```

```
{
```

```
  "output_modality": "text",
  "output_description": "The user's name as it appears in the input: {user_name}",
  "output_structure": "single item"
}
```

```
Answer paradigm: "extracted_span(s)"
```

```
// Example 3: Multiple choice selection
```

```
{
  "output_modality": "text",
  "output_description": "Choose the best summary from the list: {summary1}, {summary2}, {summary3}",
  "output_structure": "single item"
}
```

```
Answer paradigm: "one_option_from_a_set"
```

```
""
```

#### *#tasks\_and\_domains*

```
tasks_system_prompt = ""
```

You are skilled at analyzing LLM prompts and identifying their corresponding tasks and domains.

You are given a prompt text (as a raw string). Your task is to:

1. Identify the NLP/AI Task(s):

- Match the prompt to established NLP or AI task names (e.g., summarization, question answering, NLI, paraphrasing, simplification, text-generation, code-generation, code-fixing, planning, etc.).
- Use standard and general terms. Avoid overly specific labels like "medical QA" — instead, use "question answering".
- If multiple tasks apply, list them all.
- If the task cannot be identified, use "undefined".

2. Provide a Subtask for Each Task:

- Give a more granular description of what the task is doing in this case.
- For example, for task = "summarization", a possible subtask might be "article summarization".
- Every task must have a corresponding subtask.

3. Determine the Domain(s):

- Identify the domain of the prompt (e.g., medical, finance, news, legal, travel, etc.).
- Be specific and exhaustive. If unclear or unidentifiable, use "undefined".
- If multiple domains apply, list them all.

Output: Return a JSON object with exactly two fields:

- "task": a list of objects, each with:
  - "task": the standardized task name.
  - "subtask": a brief specific description.

- "domain": a list of objects, each with:
  - "domain": the domain name.

If no task or domain can be identified, use "undefined" as the value of the corresponding fields.

**\*\*Output only valid JSON.\*\***

For example:

```
{
  "task": [
    { "task": "...", "subtask": "..." },
    { "task": "...", "subtask": "..." }
  ],
  "domain": [
```

```
{ "domain": "... " },
{ "domain": "... " }
]
}
```

### #language

language\_system\_prompt = ""

You are an expert at analyzing text prompts, identifying their human languages and translating them to English.  
Note: The word "language" here refers to human languages like French, English, Chinese etc., not to programming languages!

You are given a prompt text (provided as a raw string). You have to produce a JSON object with exactly two fields:

1. "languages": a list of objects, corresponding to detected human languages in the text.  
Important! Make sure this is indeed a natural human language (like German, English etc.) and *not* a programming language\*.

Each object must include three subfields:

- "language": the language name (e.g. "English", "Chinese").
- "orig\_text": the exact substring of the prompt written in that language.
- "translated\_text": if the language is not English, its translation of that substring into English; otherwise null.

If the same language appears in separate *non-consecutive* segments, include multiple objects — one per segment — in the order they occur.

- \*Do not split *consecutive* segments in the same language\*.
- If they are consecutive they should all form one segment.
- \*Do not split if not necessary\*.

If the language of a certain span cannot be detected, return [{"language": "Undefined", "orig\_text": "<span text>", "translated\_text": None}].

Once again: **Avoid programming languages**.

2. "explicit\_language\_mentions": a list of objects for each place the prompt explicitly names a natural human language.  
Important! Make sure this is indeed a *natural human language*\* (like German, English etc.) and *not* a programming language\* or just a mention unrelated to languages (like "Turkish food").

Once again: **Avoid programming languages**.

Each object must include 2 subfields:

- "language": the named language (e.g. "French").
- "mention": the exact substring from the prompt where it is named.

- The "mention" span should include the language name and -if possible - some minimal surrounding context (for example, "respond in French").

- The "mention" span should be an exact span from the prompt.

Important: the "mention" should be an *exact span*\* from the prompt! If no human natural language name (like "English", "French", "Japanese" etc.) is used in the prompt, just return an empty list ("explicit\_language\_mentions": []).

Important: make sure *the "mention" indeed contains the language name*\* from the "language" field (for example, "French", "Spanish" etc.)

Once again: **Avoid programming languages**.

If there are no explicit mentions, return an empty list("explicit\_language\_mentions": []).

**Output only valid JSON.**

For example:

```
{
  "languages": [
    {
      "language": "Chinese",
      "orig_text": "<Text in Chinese>",
      "translated_text": "Hello, how was your day?"
    },
    {
      "language": "English",
      "orig_text": "Then share your feedback. Please, respond in French",
      "translated_text": null
    },
    {
      "language": "Chinese",
      "orig_text": "<Text in Chinese>",
      "translated_text": "Goodbye!"
    }
  ],
  "explicit_language_mentions": [
    {
      "language": "French",
      "mention": "Please respond in French"
    }
  ]
}
```

### #text extraction

text\_extraction\_system\_prompt = ""

You are an expert at extracting readable prompt text from the "messages" field of the chat.completions.create OpenAI API call parameters .

You are given the contents subfield of one of the "messages". You need to extract readable text from this "content" field if any.

- If an intelligible, readable prompt text cannot be extracted from a certain content field, return "None" (make sure you return it as a string).

- However, even if very little readable text can be extracted, you should still extract it.

- If there are variable names and/or placeholders inside the prompt text, include them in braces (or any other form which fits).

- If a variable name is enclosed into markers <VAR:...>, remove these wrapping markers but keep the variable name itself.

- Also remove code artifacts (slashes, func, args, etc.), and anything that is not part of the text as such, substituting them with a clear placeholder where applies. For example, {'func': '<VAR:image\_b64>', 'args': ['screenshot.jpg'], 'keywords': {'quality': 'high'}} should become {image\_b64(screenshot.jpg,quality=high)}.

- However, always keep any readable text even if surrounded by a lot of code remnants.

- Be sure not to omit or invent any readable text. Only render faithfully any readable text from the original message content.

- Always return your answer as a string.

""

# TalkTag: Fine-Grained Morphosyntactic Error Annotation for Transcribed Speech

Shamira Venturini<sup>1,2</sup>, Oliver Hennhöfer<sup>2</sup>, Steffen Kinkel<sup>2</sup>, Jannik Strötgen<sup>2</sup>

<sup>1</sup>Karlsruhe Institute of Technology, Germany

<sup>2</sup>Karlsruhe University of Applied Sciences, Germany

Correspondence: [shamira.venturini@h-ka.de](mailto:shamira.venturini@h-ka.de)

## Abstract

Fine-grained morphosyntactic error annotation is important in clinical and developmental language research, yet it is labour-intensive, expert-dependent, and difficult to scale. We present TalkTag, an LLM-based lightweight tool fine-tuned to automate CHAT-style error annotation in spoken-language transcripts. Developed under conditions of extreme data scarcity using children’s narrative data, the system shows the feasibility of linguistic analysis in low-resource settings. Our evaluation demonstrates that TalkTag produces encouragingly precise annotation while effectively identifying instances where linguistic ambiguity makes automated tagging genuinely complex. In summary, with TalkTag, we provide a scalable alternative to manual error annotation and practically viable support for morphosyntactic error annotation.

## 1 Introduction

Language Sample Analysis has become an increasingly important method in clinical linguistics and developmental psycholinguistics (MacWhinney and Fromm, 2022). Drawing on naturalistic spoken interaction data, it supports the study of language development and impairment in context.

During the past decades, TalkBank<sup>1</sup> (MacWhinney et al., 2004) has substantially advanced the infrastructure for this line of research through large open spoken corpora, the CHAT transcription format, and the CLAN analysis tools (MacWhinney, 2000). CHAT provides a standard way to represent spoken-language transcripts, while CLAN (Computerized Language ANalysis) provides analysis programs for CHAT-formatted files. In this setting, morphosyntactic error codes are written inline on the transcript line, immediately after the form they

describe, so that the annotation preserves both the child’s production and the analyst’s interpretation of the intended target. Although TalkBank includes tools for some aspects of automatic transcription and morphosyntactic annotation, to the best of our knowledge, it currently does not support automatic annotation of morphosyntactic errors.

However, fine-grained morphosyntactic error annotation is important because it provides evidence of grammatical development, impairment, and variation in typically developing (Moraleda-Sepúlveda and López-Resa, 2022), as well as populations such as children with developmental language disorders (Leonard and Deevy, 2020; Eadie et al., 2002; Rice and Wexler, 1996), Down syndrome (Witecy et al., 2023; Katsarou and Andreou, 2022; Penke, 2019), deaf and hard-of-hearing children with cochlear implant (Benassi et al., 2021; Golestani et al., 2018), and autism spectrum disorder (Huang and Fines-tack, 2020). Manually annotating these phenomena remains slow, labour-intensive, and dependent on expert knowledge. Unlike ordinary morphosyntactic tagging, this task often requires identifying a structured divergence between an erroneous production and an intended target, or recovering morphology that is absent from the surface string but obligatory in context. It therefore goes beyond simple sequence labelling, requiring structural linguistic reasoning rather than surface-level pattern matching.

Large Language Models (LLMs) offer a promising solution for this setup by integrating contextual modelling with schema-constrained generation (Devlin et al., 2019). Unlike assigning isolated tags token by token, LLMs can, in principle, evaluate an entire utterance holistically. This allows for the generation of well-formed inline annotations that capture syntactic dependencies and the underlying structural nature of the error.

<sup>1</sup><http://talkbank.org>

At the same time, annotated clinical and developmental language data are often scarce, access-restricted, and difficult to use for large-scale model development (Al-Marridi et al., 2026; Gagliardi and Maffia, 2024). We therefore introduce TalkTag<sup>2</sup>, a tool for fine-grained morphosyntactic error annotation that follows the CHAT guidelines for word-level error coding. The tool was developed under extremely low-resource conditions, with very limited annotated data (Hedderich et al., 2021). To address this, we employed synthetic data augmentation to fine-tune a lightweight, open-weight LLM, effectively expanding the model’s exposure to rare error patterns. In this initial prototype, we focus on children’s narrative data, a domain selected not only for its availability but for its high density of developmentally salient morphosyntactic phenomena. These narratives provide a benchmark for evaluating fine-grained annotation capabilities in a complex, real-world linguistic context.

The main contribution of this paper is a lightweight tool for fine-grained morphosyntactic error annotation in clinical and developmental language research, together with its accompanying Python package, TalkTag. More specifically, we formulate CHAT morphosyntactic error coding as a constrained structured-generation task, adapt a small instruction-tuned model to this setting under severe data sparsity, and evaluate the resulting system using automatic scoring, blinded post-hoc adjudication, and human review on unseen corpus material.

The remainder of the paper is structured as follows. Section 2 defines the target annotation scheme and the subset of CHAT morphosyntactic error labels modelled in this study. Section 3 then situates the work in relation to prior research on clinical language annotation, error coding, and automatic linguistic analysis. Section 4 describes the model, training setup, data, and evaluation design. Finally, Section 5 reports the automatic and human-reviewed results. Section 6 discusses the main linguistic error patterns, the implications of the findings for annotation practice, and the limitations of the current system.

## 2 The Annotation Language

Within the TalkBank ecosystem, which provides infrastructure for the transcription and analysis

<sup>2</sup><http://github.com/OliverHennhoefer/talk-tag>

of spoken interaction data, the MOR and GRASP programs (MacWhinney, 2012) support automatic morphosyntactic analysis of CHAT transcripts. MOR is a morphological analyser that assigns lexical and grammatical information to each token, producing the %mor tier with lemma, part-of-speech, and inflectional information. Building on this output, GRASP derives syntactic structure by assigning grammatical relations and dependency-based representations across the utterance on the %gra tier. Together, these tools make it possible to move from raw transcript text to a linguistically enriched representation of children’s speech. However, their purpose is to recover morphological and syntactic structure rather than to identify, classify, or encode morphosyntactic errors explicitly.

The CHAT Transcription Guidelines provide a general system for marking word-level errors (MacWhinney, 2000, 2019). At the level of string form, the annotation tags have a simple and regular surface structure: first, they consist of a fixed bracketed frame following the relevant error, marked with an \*. Next, a flat sequence of colon-separated fields indicates i) the error domain (phonological, semantic, neologistic, dysfluency, and morphological), ii) the error pattern (e.g., missing, superfluous, over-regularised, double-marked morphemes, unknown/known target, etc.) and iii) the morpheme or part-of-speech involved (e.g., past tense, perfective, plural, or third-person singular agreement morphemes; pronouns, prepositions, determiners as parts-of-speech). When the intended target is known, it can also be provided in brackets next to the error using the format [ : target ]. This is used when CLAN’s MOR morphological analyser should analyse the target form instead of the produced one, whereas [ : : target ]<sup>3</sup> can be used to preserve analysis of the produced real-word form while still recording the intended target.

The surface syntax of the labels is therefore relatively simple, yet their assignment is highly challenging: it often depends on contextual linguistic interpretation and, in many cases, on reconstructing an intended target form.

In this work, we focus on treating [\* m : \*] labels as mismatches between produced and target forms

<sup>3</sup>Since the time of this study, TalkBank has updated the CHAT manual, replacing the [ : : target ] syntax with [= target ]. While the model described here was trained on the earlier convention, the associated Python package includes post-processing to ensure compliance with the latest standard and offers options to toggle the visibility of reconstructions depending on user preference.



Figure 1: The TalkTag Workflow Pipeline.

under lexical identity (i.e., morpheme operations), and  $[* s:r:*$ ] labels as substitutions involving the same lexical category (e.g., wrong preposition) or  $[* s:r:gc:*$ ] wrong grammatical category (e.g., adjective for pronoun). Moreover, we use both reconstruction strategies, reserving  $[: target]$  for cases where the error produces a non-word form. Examples of morphological error annotations are:

- "*Yesterday I walk*  $[:: walked]$   $[* m:\emptyset ed]$  *to school*", which marks a missing past tense morpheme.
- "*Yesterday I goed*  $[: went]$   $[* m:=ed]$  *to school*", which marks an overregularised past tense morpheme resulting in a non-word form.

Examples of substitutional error annotations are:

- "*Yesterday me*  $[:: I]$   $[* s:r:gc:pro]$  *walked to school*", which marks a wrong grammatical category of a pronoun.
- "*Yesterday I went in*  $[:: to]$   $[* s:r:prep]$  *school*", which marks a wrong preposition.

The CHAT manual provides a standard inventory of error codes, but the framework is extendable: CHAT coding can be adapted to specific applications, and the error-coding system itself allows additional distinctions and combinations within that general format. The annotation scheme for morphological and substitutional errors is illustrated in Table 1 and Table 2, respectively. The complete label components inventory is illustrated in Table 11 in Appendix D.

### 3 Related Work

**Grammatical error detection.** Work on grammatical error detection and grammatical error correction addresses linguistic errors more directly, but typically formulates the problem as one of edit detection or sentence-level correction. In this literature, annotation generally starts from a source sentence and a corrected target, from which error

Level 1	Meaning
* m:	morphosyntactic error
Level 2	Meaning
$\emptyset$	missing regular form
base:	base for irregular form
irr:	irregular for base form
sub:	past/perfective substitution
=	over-regularisation
+	superfluous marking
++	double marking
vsg:	irregular verb 3PS
vun:	irregular verb unmarked
allo	allomorphic errors
Level 3	Meaning
a	agreement error
i	irregular target
mor	target morpheme

Table 1: CHAT annotation scheme for morphological errors.

spans are identified and labelled with edit operations such as replacement, omission, insertion, or transposition. Schemes such as ERRANT (Korre and Pavlopoulos, 2020) add a further layer of linguistic classification, yielding a structured representation of each edit rather than a flat label. For example, ERRANT combines edit operations such as M (missing), R (replacement), and U (unnecessary) with linguistic error categories to produce fine-grained composite labels. However, the empirical basis of this literature is overwhelmingly written and learner-focused. As noted by Bryant et al. (2023), the main benchmark datasets are largely derived from L2 English essays, examinations, and learner-platform submissions, including FCE (Yannakoudakis et al., 2011), NUCLE (Dahlmeier et al., 2013), CoNLL-2013 and 2014 (Ng et al., 2013, 2014), Lang-8 (Mizumoto et al., 2012; Tajiri et al., 2012), JFLEG (Napoles et al., 2017), and W&I+LOCNESS (Bryant et al., 2019). Consequently, current annotation schemes and correction models are primarily optimized for

Level 1	Meaning
* s :	substitution error
Level 2	Meaning
r :	related lexical substitution
r :gc :	related grammatical substitution
Level 3	Meaning
POS	target morpheme or part-of-speech

Table 2: CHAT annotation scheme for substitutional errors.

sentence-level written L2 language, which limits direct transfer to spoken, interactional, or clinically atypical language.

**Morphosyntactic error annotation of child language.** More specific work on child language appears limited and idiosyncratic in terms of the annotation scheme and targeted granularity. [Morley et al. \(2013\)](#) first showed that even relatively coarse linguistic error codes could be sufficient for identifying neurodevelopmental disorders. Building on that result, [Morley et al. \(2014\)](#) developed a data-driven dependency-parser approach for detecting and labelling grammatical errors in SALT-annotated transcripts of children’s speech. SALT ([Miller et al., 2011](#)) supports error coding through a relatively small default inventory of labels, including over-generalization [EO: ], pronoun error [EP: ], other word-level error [EW: ], extraneous word [EW], and utterance-level error [EU], although the active code set can be customised within the software. [Morley et al. \(2013\)](#) evaluated on the ENNI corpus from CHILDES and the NSR corpus from the SALT database, their system outperformed both Microsoft Word’s grammar checker and a Naive Bayes baseline, while also showing that performance was sensitive to corpus-specific annotation practices and differences in label granularity. This work is therefore highly relevant to the present study, but it remains grounded in a relatively coarse inventory rather than a more fine-grained annotation scheme.

Earlier work by [Hassanali and Liu \(2011\)](#) explored a more fine-grained approach to grammatical error annotation in child language transcripts. Using 677 transcripts from the Paradise corpus ([Paradise et al., 2005](#)), they manually annotated ten error categories, with particular attention to verb-related errors, and compared rule-based parse-template methods with statistical classifiers for de-

tecting six error types. Their results showed that statistical approaches generally outperformed rule-based ones. The study’s main contribution was to show that automatic grammar checking could move beyond holistic measures of syntactic development, such as IPSyn ([Sagae et al., 2005](#)), providing a more differentiated profile of grammatical weaknesses. At the same time, it highlighted important limitations, including the difficulty of parsing spoken child language with disfluencies and incomplete utterances, ambiguity in assigning error categories, and the restricted coverage of systems built around a narrow set of constructions.

More recent work has also approached child grammar from other angles: [Nikolaus et al. \(2024\)](#) developed a context-sensitive scheme for annotating child utterances in caregiver conversation as grammatical, ungrammatical, or ambiguous, and trained Transformer-based models on 4,200 manually annotated CHILDES utterances. Their best models reached near-human agreement and were used to annotate a much larger corpus, confirming that grammaticality increases with age. Unlike work targeting explicit morphosyntactic error labelling, however, their focus was on broad utterance-level grammaticality in a conversational context.

Most recently, [Gebauer et al. \(2025\)](#) investigated grammatical error detection in spontaneous children’s speech using German kidsTALC data ([Rumberg et al., 2022](#)), explicitly addressing both ASR errors and ambiguity in manual error labelling. They proposed a BERT-based recurrent model with iterative pseudo-labelling, showing significant improvements on both manual and automatically transcribed speech. This makes their study particularly relevant to the present work, since it tackles realistic spoken child-language data rather than written text alone. However, the task is still formulated as coarse binary error detection rather than fine-grained morphosyntactic error annotation, so it provides a close methodological precedent without addressing the richer annotation language used in this paper.

Taken together, the literature points to a clear research gap. While grammatical error classification and correction are well developed, this work is largely grounded in written learner-language data and does not transfer straightforwardly to clinical or developmental spoken-language settings. The more specific literature on child spoken-language error analysis is comparatively sparse, and the clos-

est prior systems either date back more than a decade or adopt different annotation scopes and levels of granularity. To our knowledge, there is currently no automated tool for fine-grained, CHAT-compatible morphosyntactic error annotation within the TalkBank ecosystem.

## 4 Methods

This section describes the study’s methodological setup: the model architecture, training regime, data sources, and evaluation strategy for the morphosyntactic error annotation. The workflow pipeline is visually illustrated in Figure 1.

**Model.** The model is instruction-tuned base Meta-Llama-3.1-8B-Instruct (Grattafiori et al., 2024), loaded in 4-bit quantised form (bnb-4bit) for efficient fine-tuning. Training is implemented using the Unsloth framework, enabling parameter-efficient adaptation via LoRA while keeping the base model weights frozen. This setup allows us to fine-tune an 8B-parameter model under constrained hardware conditions.

The model is instruction-tuned to produce exactly one annotated utterance line, while preserving the original token order, spelling, punctuation, disfluencies, and CHAT symbols. The prompt constrains the model by specifying the structural conditions of the annotation language. The full prompt is provided in Appendix A.

Rather than treating CHAT morphosyntactic labels as a flat inventory of opaque output strings, we treat the annotation scheme as a structured symbolic language. Accordingly, we extend the tokeniser not with full label forms, but with reusable components that recur across the annotation system, including bracketed markers, domain indicators, and subtype fragments. The embedding matrix is resized to accommodate this augmented vocabulary. This reduces fragmentation of CHAT-specific sequences under the base tokeniser and encourages the model to compose licensed tags from meaningful subparts rather than retrieve them from a closed set. This design is motivated by the structure of the annotation scheme itself: CHAT error tags are compositionally organised, encoding contrasts such as domain, operation type, agreement sensitivity, and irregular morphology. The model is thus exposed to the building blocks of a small regular annotation language and must learn to generate well-formed tag combinations under task constraints.

Source	Total	0	1	2	$\geq 3$
Synthetic	5830	946	4771	109	4
Real	4585	4015	517	47	6
Total	10415	4961	5288	156	10

Table 3: Distribution of utterances by number of annotated errors in the full pre-split master dataset. Columns 0, 1, 2, and  $\geq 3$  indicate the number of utterances containing zero, one, two, or three or more errors, respectively.

**Training.** Fine-tuning uses LoRA with rank 32,  $\alpha = 64$ , and dropout 0.05. Maximum sequence length is 384 tokens. Training uses a per-device batch size of 8 with gradient accumulation of 4, a warmup ratio of 0.03, and weight decay of 0.1. The model was fine-tuned on an NVIDIA A100-SXM4-80GB GPU. Training took approximately 47 minutes for 3 epochs and 723 optimiser steps.

**Data.** We use CHAT-formatted utterances drawn from a subset of the Edmonton Narrative Norms Instrument (ENNI) corpus (Schneider et al., 2006), available through TalkBank/CHILDES<sup>4</sup> (MacWhinney, 2000). We focus on narratives from children aged 4-5, since this developmental range provides a significant concentration of morphosyntactic phenomena, including overregularisation, agreement errors, tense marking, and clause linking (Cummings, 2023). The resulting real-data subset contains 4,585 utterances manually reviewed for the target annotation task. As shown in Table 3, this is an extremely low-resource setting (Hedderich et al., 2021) not only in corpus size but also in error density, since most real utterances are error-free and many labels have very limited support.

To mitigate this sparsity, we supplemented the real corpus with curated synthetic examples covering configurations described in the CHAT guidelines. These were generated from error-conditioned prompts and manually reviewed before inclusion. We targeted a minimum overall support of approximately 100 instances per label. We also retained clean utterances in both the real and synthetic portions of the data so that the model would learn when to leave an utterance unchanged and preserve valid CHAT formatting.

**Evaluation.** We evaluate the system at three complementary levels: automatic scoring on held-out

<sup>4</sup><https://talkbank.org/childes>

data, followed by a blinded post-hoc review of disagreement cases, and human evaluation of unseen data annotation.

For automatic evaluation, we split both the real ENNI data into train and validation sets as the primary confirmatory benchmarks, as well as augmented data to be used as label-coverage diagnostics. To stabilise rare-label evaluation, we enforced minimum per-label support in the synthetic coverage splits of  $N = 10$ , without distorting the natural error distribution of the real set. This design allows rare-label behaviour to be measured under controlled support while preserving the full real training pool and avoiding aggressive downsampling.

Automatic evaluation combines line-level and label-level perspectives. We report exact match over the full annotated utterance line, but treat it as a secondary summary measure, since the high proportion of clean utterances inflates this score. Our main evaluation metrics focus on the error labels themselves: micro-F1, macro-F1, and per-label precision/recall/F1. Per-label results are further divided into confirmatory and exploratory subsets using minimum-support thresholds, so that claims about individual labels are not based on extremely small counts. Reconstruction targets are reported, but they are not treated as the primary object of evaluation.

Since manual reference annotation may be incomplete or underspecified, automatic scoring can over-penalise outputs that are linguistically plausible but do not match the gold line exactly. We therefore conduct blinded post-hoc adjudication on disagreement cases from the final model. For each reviewed utterance, the reviewer sees the original input and candidate annotations without knowing whether a candidate comes from the gold reference or from the model. Each candidate is assigned one of four labels (correct, incorrect, ambiguous, or unsure) and particularly informative cases are flagged for qualitative analysis. Source identity is stored separately and revealed only after review is complete. These judgments are then merged with the hidden source labels to estimate how often apparent automatic errors reflect genuine model failures, as opposed to ambiguity or omission in the reference annotation. We finally report a conservative updated exact match estimate for the label-bearing subset.

To assess generalisation to unseen data from the same corpus, we run the final model on the re-

mainder of the ENNI corpus, comprising 13,637 utterances not used for training or in-domain evaluation. We manually review all 854 utterances for which the model produced an error annotation, a random sample of 2,200 clean unannotated utterances to estimate the frequency of missed errors, and 91 unannotated but modified utterances. The total reviewed sample is therefore 3,145 utterances. We apply the same coding labels as in the post-hoc review of the automatic evaluation.

## 5 Results

This section reports the results of TalkTag across three complementary stages of evaluation: automatic scoring on held-out training data, blinded post-hoc adjudication, and human review on test inference data. The main question is not whether the system can replace expert annotation, but whether it can provide a reliable first pass that surfaces plausible morphosyntactic error candidates for review.

### 5.1 Automatic Model Evaluation

As shown in Table 4, on the primary test split, the final model achieves 93.6% exact match, 86.0% micro-precision, 75.5% micro-recall, and 80.4% micro-F1. Validation performance is slightly higher, with 95.4% exact match and 89.0% micro-F1. On the synthetic support split, the model reaches 93.2% micro-F1. Because the real test split is dominated by clean utterances (600/687), we also report results restricted to the 87 utterances containing at least one gold error label. On this subset, the model achieves 66.7% exact match, 91.4% precision, 75.5% recall, and 82.7% micro-F1. The strongest per-label results on the real test data are obtained for overregularised past morphology (F1 89.4,  $N = 26$ ), missing third-person singular marking (F1 86.3,  $N = 25$ ), and allomorphic errors (F1 100.0,  $N = 14$ ), while pronoun substitution is lower at F1 77.8 ( $N = 10$ ). This pattern suggests a high-precision annotation aid: proposed labels are usually plausible, but the lower recall indicates that some linguistically valid errors remain difficult to recover automatically, especially when the intended target depends on wider discourse context.

### 5.2 Post-hoc Review

Aware that the manual gold annotations may themselves contain errors or omissions, we carried out blinded post-hoc adjudication on all 44 disagreement utterances in the test set: 20 reviewed disagreements were judged acceptable for both model

Split	N	EM	P	R	F1
Val.	687	95.4%	92.0%	86.2%	89.0%
Test	687	93.6%	86.0%	75.5%	80.4%
Labels	87	66.7%	91.4%	75.5%	<b>82.7%</b>
Post-hoc	87	<b>82.8%</b>			

Table 4: Automatic evaluation results. EM denotes full-line exact match; P, R, and F1 are micro-averaged over error tags. Labels restrict evaluation to the 87 tagged utterances in Test; Post-hoc refers to the increased score after manual review of disagreement cases.

and reference, and in 5 further cases the model output was preferred, yielding a conservative post-hoc acceptable-output rate of 82.8% (72/87) error-bearing utterances alone.

All the reported results on the real test set are displayed in Table 4, while synthetic support split (Table 7), label-wise results (Table 9), and post-hoc details (Table 8) can be found in the Appendix.

### 5.3 Human Evaluation of Model Outputs

Results of the model’s annotation of the unseen portion of the ENNI corpus are summarised in Table 5. At the utterance level, 93.4% of reviewed outputs were judged acceptable and 6.6% incorrect. At the label level, 83.7% of reviewed label-bearing cases were judged correct. Among the 146 unaccepted label judgments, 94 were false positives, 31 were incorrect labels, and 20 were false negatives. In the audit of 2,200 model-clean utterances, 7 probable missed errors were identified.

At the label level, agreement-related labels were the most frequent in the reviewed predictions. Grouping the three main agreement labels – missing third-person singular marking (e.g., *he go* for *goes*), irregular unmarked for singular (e.g., *he are/were* for *is/was*), and irregular singular for unmarked (e.g., *they is/was* for *are/were*) – yields 408 reviewed instances in total. The next most frequent labels were allomorphic errors (105), overregularised past forms (84), and pronoun substitutions (70). Among past-related labels, the most frequent pattern was substitution of the base form for an irregular past form (e.g., *go* for *went*). Accuracy for these frequent labels was 87.3% for the agreement group as a whole, 92.4% for allomorphic errors, 78.6% for overregularised past forms, 71.4% for pronoun substitutions, and 82.1% for base-for-irregular past substitutions. Detailed per-label results are reported in Table 10 in the Appendix.

Measure	Count	Percent
Reviewed utterances (raw)	3145	100.0%
Out-of-scope exclusions	22	0.7%
Official reviewed total	3123	–
Utterance-level acceptable	2917	93.4%
Utterance-level incorrect	206	6.6%
Label-level reviewed cases	894	–
Label-level correct	748	83.7%
Label-level incorrect	146	16.3%
Incorrect label	31	3.5%
False negative	20	2.2%
False positive	94	10.5%
Audited model-clean utterances	2200	–
Probable missed errors	7	0.32%

Table 5: Human-reviewed evaluation on unseen ENNI data. Percentages for utterance-level outcomes are computed over the reviewed total ( $N = 3123$ ). Label-level percentages are computed over reviewed label-bearing cases ( $N = 894$ ). The clean-audit miss rate is computed over the audited model-clean sample ( $N = 2200$ ).

Among the qualitatively reviewed cases, the clearest and most recurrent linguistic patterns causing confusion involved uninflected verb forms, especially the interaction between tense and agreement and cases of invariant verbs (irregular verbs whose past form is identical to the base form). We focus on this pattern here because it directly addresses a marker of developmental language-disordered speech: preference for uninflected verb forms.

We found 16/17 cases in which the model erroneously assigned the agreement label to uninflected verbs following a third-person singular subject, even though an obligatory context for licensing the past tense was present. Of these, 8 cases involved an invariant verb. We found five additional cases in which invariant verbs were overtly overregularised (e.g., *hurted*, *costed*, *putted*); in these cases, the model did not converge on a single analysis but alternated between overregularisation and double-marking annotation.

## 6 Discussion

Taken together, the results suggest that a useful tool for fine-grained morphosyntactic error annotation can be developed even under conditions of extreme data sparsity and limited computational resources. Although TalkTag’s scores on the full test set are partly inflated by the large number of error-free utterances, performance remains encouraging on genuinely error-bearing cases. In addition, the blinded post-hoc review shows that a non-trivial subset of apparent disagreements reflects linguistic

ambiguity or underspecification in the reference annotation rather than straightforward model failure. This is important both for metric interpretation and for practical use: full-line disagreement might not always correspond to a linguistically unacceptable output.

**Clinical and developmental relevance.** The results are especially encouraging for agreement-related errors, overregularised past forms and pronoun substitutions, which were among the most frequent and best-supported categories in the reviewed data. These categories are also linguistically meaningful. Tense and agreement morphology, as well as difficulties with pronouns, are central to the study of developmental language disorder and autism spectrum disorder (Leonard and Deevy, 2020; Eadie et al., 2002; Rice et al., 1998; Wexler et al., 1998; Rice and Wexler, 1996). Overregularisation is a well-established feature of typical language development and remains informative when it persists or occurs at elevated rates in atypical development (Moraleda-Sepúlveda and López-Resa, 2022; Marcus et al., 1992). From this perspective, the model’s relative success on these labels is encouraging not only in engineering terms but also because it aligns with clinically and developmentally relevant dimensions of child language.

**Agreement bias in tense-agreement ambiguities.** At the same time, the qualitative review revealed a clear and recurrent failure mode involving bare verb forms following third-person singular subjects. In these cases, the model often preferred agreement-based analyses over missing past-tense interpretations, including in some contexts where the surrounding discourse licensed a past reading. This pattern was especially pronounced for zero-change verbs such as *hurt* and *put*, whose past forms are identical to their base forms. More generally, zero-change verbs formed a persistent challenge in the reviewed sample, and when they were overtly overregularised (e.g., *hurted*, *costed*, *putted*), the model alternated between overregularisation and double-marking analyses rather than converging on a single label decision.

These patterns are informative because they suggest that the model’s errors are not simply random, but partly structured by a preference for locally recoverable agreement analyses over broader tense interpretation. One plausible explanation is that agreement errors are both more locally diagnosable and more strongly represented in the training

signal, whereas missing past-tense interpretations often require integration of wider temporal and discourse context. At the same time, these cases also highlight a genuine property of the task itself: in spontaneous non-standard speech, morphosyntactic annotation is often difficult precisely because sparse morphology, discourse context, and lexical irregularity interact.

**Generalisation and pre-annotation utility.** The broader review on unseen ENNI material seems to reinforce this picture. Although overall acceptability remained high, the error profile on this larger sample differed somewhat from that of the held-out test set: agreement-related labels again dominated the reviewed predictions, but the main source of degradation was over-annotation rather than omission. The audit of model-clean utterances nevertheless suggests that silent misses remain comparatively infrequent. This pattern supports a cautious interpretation of the tool as one that is more useful for surfacing plausible candidate errors than for producing final annotations without review.

From a practical perspective, the present results suggest that the tool is already useful as a pre-annotation aid, even where its outputs still require human correction. Prior work on machine-assisted annotation has shown that automatic pre-annotation can reduce annotation effort, improve consistency, and increase annotation speed without harming final quality (Lingren et al., 2014; Fort and Sagot, 2010; Mikulová et al., 2022). The contribution of the present system is therefore not that it eliminates the need for expert review, but that it provides a linguistically informed first-pass annotation over a large volume of CHAT material. This is especially valuable in a domain where fully manual annotation is slow, costly, and itself subject to ambiguity and inconsistency.

**Scope of the present prototype.** At the same time, the present findings should be interpreted within the scope of the current prototype. The model was developed and evaluated on children’s narrative data from a single corpus family, and the synthetic support split serves only as secondary evidence of label-space coverage under controlled conditions rather than as a substitute for naturalistic evaluation. Moreover, some reviewed errors were tied to pre-existing CHAT annotations or discourse-formatting cues rather than to morphosyntactic analysis alone.

## Limitations

The present study should be understood as a prototype rather than a complete solution to CHAT-style error annotation. The model was developed under conditions of extreme data scarcity and trained on a reduced subset of the CHAT error inventory, focusing on selected morphosyntactic and closely related substitution labels. As a result, the current system does not yet cover the full range of CHAT-compatible error phenomena, and support for some rare or irregular patterns remains limited. This is particularly relevant for error types that were only sparsely represented in the available training data. Importantly, this is not simply an artefact of the present experiment: some of these phenomena are genuinely infrequent in naturalistic corpora, which makes it inherently difficult to obtain enough real examples for robust learning.

The current annotation scope is also intentionally narrow. For example, forms such as *they going to the shop* may plausibly be analysed not as cases of superfluous progressive marking, but as instances of missing auxiliary. However, the present model was not trained to annotate missing parts of speech, and such cases therefore fall outside the annotation scope of the current system. Within that restricted label space, we treated the model's superfluous-progressive analysis as acceptable, since it captures a genuine deviation while avoiding unsupported labels. More generally, the prototype is better understood as a first step toward fine-grained morphosyntactic error annotation than as a holistic grammar annotation tool.

These constraints also make it important to understand what the model is learning under conditions of scarcity. The present study does not disentangle the mechanisms by which the model produces its annotations. The fine-tuning was designed to encourage recovery of an intended target form while generating the final inline label, making the approach loosely related to an analysis-by-synthesis perspective. At the same time, the model was also allowed to generate CHAT labels compositionally rather than retrieve them from a fixed inventory of whole forms. The current results, therefore, do not establish whether the observed gains arise from implicit target-form reconstruction, from better formatting control, or from a learned mapping between linguistic error patterns and the internal structure of the annotation language. Future work will investigate these possibilities more directly, in-

cluding whether improvements are concentrated in error types that require target-form recovery and whether performance depends on latent correction-like inference.

A further limitation concerns generalisation. Although the model was evaluated both on held-out data and on new unseen material, all real-data evaluation was conducted within the ENNI corpus. This provides a useful first test of robustness, but it does not establish how well the model transfers across other CHILDES corpora, elicitation settings, age ranges, or clinical populations. Future work should therefore test the system on additional child-language corpora, as well as on more clearly out-of-distribution material such as second-language learner data and adult's clinical speech.

These limitations point to the main conditions under which the prototype should be used: as a human-in-the-loop pre-annotation aid rather than as a fully automatic replacement for expert judgement. Broader annotation coverage, greater robustness to authentic CHAT markup, and wider evaluation across corpora and populations will be needed before the system can be treated as a more general annotation tool.

## Acknowledgments

This work was supported by the Ministry of Science, Research and the Arts of Baden-Württemberg (MWK) and KIT's Accessibility through AI-based Assistive Technology (KATE) Graduate School.

**Generative AI Assistance Declaration** During the preparation of this work, the author(s) used ChatGPT to rephrase, proofread or summarise text content. After using this tool, the author(s) reviewed and edited the content as needed and take(s) full responsibility for the publication's content.

**Data Availability Statement** The data used in this study are not redistributed with our code release. They are hosted by TalkBank/CHILDES and should be obtained directly from the official source under the applicable TalkBank access and licensing rules.

## References

- Abeer Z. Al-Marridi, Samawiyah M. Ulde, Ahmed Bensaid, and Tariq A. Khwaileh. 2026. [Speech and language disorders: A systematic review of corpora and future directions](#). *Applied Corpus Linguistics*, 6(1):100186.

- Erika Benassi, Sonia Boria, Maria Teresa Berghenti, Michela Camia, Maristella Scorza, and Giuseppe Cossu. 2021. [Morpho-syntactic deficit in children with cochlear implant: Consequence of hearing loss or concomitant impairment to the language system?](#) *International Journal of Environmental Research and Public Health*, 18(18):9475.
- Christopher Bryant, Mariano Felice, Øistein E. Andersen, and Ted Briscoe. 2019. [The BEA-2019 shared task on grammatical error correction](#). In *Proceedings of the Fourteenth Workshop on Innovative Use of NLP for Building Educational Applications*, pages 52–75, Florence, Italy. Association for Computational Linguistics.
- Christopher Bryant, Zheng Yuan, Muhammad Reza Qorib, Hannan Cao, Hwee Tou Ng, and Ted Briscoe. 2023. [Grammatical error correction: A survey of the state of the art](#). *Computational Linguistics*, pages 1–59.
- Louise Cummings. 2023. [Communication disorders: A complex population in healthcare](#). *Language and Health*, 1(2):12–19.
- Daniel Dahlmeier, Hwee Tou Ng, and Siew Mei Wu. 2013. [Building a large annotated corpus of learner English: The NUS corpus of learner English](#). In *Proceedings of the Eighth Workshop on Innovative Use of NLP for Building Educational Applications*, pages 22–31, Atlanta, Georgia. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- P. A. Eadie, M. E. Fey, J. M. Douglas, and C. L. Parsons. 2002. [Profiles of grammatical morphology and sentence imitation in children with specific language impairment and down syndrome](#). *Journal of Speech, Language, and Hearing Research*, 45(4):720–732.
- Karën Fort and Benoît Sagot. 2010. [Influence of pre-annotation on POS-tagged corpus development](#). In *Proceedings of the Fourth Linguistic Annotation Workshop*, pages 56–63, Uppsala, Sweden. Association for Computational Linguistics.
- Gloria Gagliardi and Marta Maffia. 2024. [Language resources for clinical linguistics: introduction to the special issue](#). *Language Resources and Evaluation*, 58(3):859–863.
- Christopher Gebauer, Lars Rumberg, Lars Köhn, Hanna Ehlert, Edith Beaulac, and Jörn Ostermann. 2025. [Grammatical Error Detection on Spontaneous Children’s Speech Using Iterative Pseudo Labeling](#). In *Interspeech 2025*, pages 2865–2869.
- Samane Dehghani Golestani, Nahid Jalilevand, and Mohammad Kamali. 2018. [A comparison of morpho-syntactic abilities in deaf children with cochlear implant and 5-year-old normal-hearing children](#). *International Journal of Pediatric Otorhinolaryngology*, 110:27–30.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, Amy Yang, Angela Fan, Anirudh Goyal, Anthony Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev, Arthur Hinsvark, and 542 others. 2024. [The llama 3 herd of models](#). *arXiv preprint*.
- Khairun-nisa Hassanali and Yang Liu. 2011. [Measuring language development in early childhood education: A case study of grammar checking in child language transcripts](#). In *Proceedings of the Sixth Workshop on Innovative Use of NLP for Building Educational Applications*, pages 87–95, Portland, Oregon. Association for Computational Linguistics.
- Michael A. Hedderich, Lukas Lange, Heike Adel, Jan-nik Strötgen, and Dietrich Klakow. 2021. [A survey on recent approaches for natural language processing in low-resource scenarios](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2545–2568, Online. Association for Computational Linguistics.
- Timothy Huang and Lizbeth Finestack. 2020. [Comparing morphosyntactic profiles of children with developmental language disorder or language disorder associated with autism spectrum disorder](#). *American Journal of Speech-Language Pathology*, 29(2):714–731.
- Dimitra Katsarou and Georgia Andreou. 2022. [Morphosyntactic abilities in young children with down syndrome: Evidence from the greek language](#). *International Journal of Language & Communication Disorders*, 57(5):937–947.
- Katerina Korre and John Pavlopoulos. 2020. [ERRANT: Assessing and improving grammatical error type classification](#). In *Proceedings of the 4th Joint SIGHUM Workshop on Computational Linguistics for Cultural Heritage, Social Sciences, Humanities and Literature*, pages 85–89, Online. International Committee on Computational Linguistics.
- Laurence B. Leonard and Patricia Deevy. 2020. [Retrieval practice and word learning in children with specific language impairment and their typically developing peers](#). *Journal of Speech, Language, and Hearing Research*, 63(10):3252–3262.
- Todd Lingren, Louise Deleger, Katalin Molnar, Haijun Zhai, Jareen Meizen-Derr, Megan Kaiser, Laura Stoutenborough, Qi Li, and Imre Solti. 2014. [Evaluating the impact of pre-annotation on annotation speed and potential bias: natural language processing](#)

- gold standard development for clinical named entity recognition in clinical trial announcements. *Journal of the American Medical Informatics Association*, 21(3):406–413.
- Brian MacWhinney. 2000. *The CHILDES project: The database*, volume 2. Psychology Press.
- Brian MacWhinney. 2012. Morphosyntactic analysis of the CHILDES and TalkBank corpora. In *Proceedings of the Eighth International Conference on Language Resources and Evaluation (LREC'12)*, pages 2375–2380, Istanbul, Turkey. European Language Resources Association (ELRA).
- Brian MacWhinney. 2019. *Chat manual*.
- Brian MacWhinney, Steven Bird, Christopher Cieri, and Craig Martell. 2004. Talkbank: Building an open unified multimodal database of communicative interaction. In *Proceedings of the Fourth International Conference on Language Resources and Evaluation (LREC'04)*, Lisbon, Portugal. European Language Resources Association (ELRA).
- Brian MacWhinney and Davida Fromm. 2022. Language sample analysis with talkbank: An update and review. *Frontiers in Communication*, 7.
- Gary F. Marcus, Steven Pinker, Michael Ullman, Michelle Hollander, T. John Rosen, and Fei Xu. 1992. Overregularization in language acquisition. *Monographs of the Society for Research in Child Development*, 57(4):1–182.
- Marie Mikulová, Milan Straka, Jan Štěpánek, Barbora Štěpánková, and Jan Hajic. 2022. Quality and efficiency of manual annotation: Pre-annotation bias. In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 2909–2918, Marseille, France. European Language Resources Association.
- Jon F. Miller, Karen Andriacchi, and Ann Nockerts. 2011. *Assessing Language Production Using SALT Software: A Clinician's Guide to Language Sample Analysis*. SALT Software, LLC.
- Tomoya Mizumoto, Yuta Hayashibe, Mamoru Komachi, Masaaki Nagata, and Yuji Matsumoto. 2012. The effect of learner corpus size in grammatical error correction of ESL writings. In *Proceedings of COLING 2012: Posters*, pages 863–872, Mumbai, India. The COLING 2012 Organizing Committee.
- Esther Moraleda-Sepúlveda and Patricia López-Resa. 2022. Morphological difficulties in people with developmental language disorder. *Children*, 9(2):125.
- Eric Morley, Anna Eva Hallin, and Brian Roark. 2014. Data driven grammatical error detection in transcripts of children's speech. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 980–989, Doha, Qatar. Association for Computational Linguistics.
- Eric Morley, Brian Roark, and Jan P. H. Santen. 2013. The utility of manual and automatic linguistic error codes for identifying neurodevelopmental disorders. In *BEA@NAACL-HLT*.
- Courtney Napoles, Keisuke Sakaguchi, and Joel Tetreault. 2017. JFLEG: A fluency corpus and benchmark for grammatical error correction. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pages 229–234, Valencia, Spain. Association for Computational Linguistics.
- Hwee Tou Ng, Siew Mei Wu, Ted Briscoe, Christian Hadiwinoto, Raymond Hendy Susanto, and Christopher Bryant. 2014. The CoNLL-2014 shared task on grammatical error correction. In *Proceedings of the Eighteenth Conference on Computational Natural Language Learning: Shared Task*, pages 1–14, Baltimore, Maryland. Association for Computational Linguistics.
- Hwee Tou Ng, Siew Mei Wu, Yuanbin Wu, Christian Hadiwinoto, and Joel Tetreault. 2013. The CoNLL-2013 shared task on grammatical error correction. In *Proceedings of the Seventeenth Conference on Computational Natural Language Learning: Shared Task*, pages 1–12, Sofia, Bulgaria. Association for Computational Linguistics.
- Mitja Nikolaus, Abhishek Agrawal, Petros Kaklamanis, Alex Warstadt, and Abdellah Fourtassi. 2024. Automatic annotation of grammaticality in child-caregiver conversations. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 1832–1844, Torino, Italia. ELRA and ICCL.
- Jack L. Paradise, Thomas F. Campbell, Christine A. Dollaghan, Heidi M. Feldman, Beverly S. Bernard, D. Kathleen Colborn, Howard E. Rockette, Janine E. Janosky, Dayna L. Pitcairn, Marcia Kurs-Lasky, Diane L. Sabo, and Clyde G. Smith. 2005. Developmental outcomes after early or delayed insertion of tympanostomy tubes. *New England Journal of Medicine*, 353(6):576–586.
- Martina Penke. 2019. Regular and irregular inflection in down syndrome – new evidence from german. *Cortex*, 116:192–208.
- Mabel L. Rice and Kenneth Wexler. 1996. Toward tense as a clinical marker of specific language impairment in english-speaking children. *Journal of Speech and Hearing Research*, 39(6):1239–1257.
- Mabel L. Rice, Kenneth Wexler, and Scott Hershberger. 1998. Tense over time: The longitudinal course of tense acquisition in children with specific language impairment. *Journal of Speech, Language, and Hearing Research*, 41(6):1412–1431.
- Lars Rumberg, Christopher Gebauer, Hanna Ehlert, Maren Wallbaum, Lena Bornholt, Jörn Ostermann, and Ulrike Lüdtke. 2022. *kidstalc: A corpus of*

- 3- to 11-year-old german children's connected natural speech. In *Interspeech 2022*, interspeech\_2022, pages 5160–5164. ISCA.
- Kenji Sagae, Alon Lavie, and Brian MacWhinney. 2005. Automatic measurement of syntactic development in child language. In *Proceedings of the 43rd Annual Meeting of the Association for Computational Linguistics (ACL'05)*, pages 197–204, Ann Arbor, Michigan. Association for Computational Linguistics.
- Phyllis Schneider, Denyse Hayward, and Rita Vis Dubé. 2006. Edmonton narrative norms instrument.
- Toshikazu Tajiri, Mamoru Komachi, and Yuji Matsumoto. 2012. Tense and aspect error correction for ESL learners using global context. In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 198–202, Jeju Island, Korea. Association for Computational Linguistics.
- Kenneth Wexler, Carson T. Schütze, and Mabel Rice. 1998. Subject case in children with SLI and unaffected controls: Evidence for the Agr/Tns omission model. *Language Acquisition*, 7(2–4):317–344.
- Bernadette Wittecy, Eva Wimmer, Isabel Neitzel, and Martina Penke. 2023. Morphosyntactic development in german-speaking individuals with down syndrome—longitudinal data. *Frontiers in Psychology*, 14.
- Helen Yannakoudakis, Ted Briscoe, and Ben Medlock. 2011. A new dataset and method for automatically grading ESOL texts. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 180–189, Portland, Oregon, USA. Association for Computational Linguistics.

## A Annotation Prompt

**Role.** You are a TalkBank CHAT annotator for morphosyntactic error coding.

**Task.** Annotate the input utterance by inserting valid CHAT error tags inline.

**Output requirements.**

1. Preserve original token order, spelling, casing, punctuation, disfluencies, and CHAT symbols.
2. Do not rewrite, paraphrase, or correct the utterance.
3. Insert only error tags inline, following the error token.
4. If no target error is present, return the utterance unchanged.
5. Write the correct target form as [ : target] when the incorrect morpheme yields a nonword, and as [ : : target] when the error is an attested word.
6. Build each CHAT error tag compositionally from licensed scheme parts rather than relying on a memorised whole-label form.
7. Use m:\* only for same-lexeme morphological contrasts and s:\* only for substitutional contrasts.
8. Use :a only for agreement-sensitive labels that license it.
9. Use :i only where an irregular-sensitive label licenses it.
10. Output only licensed CHAT tags; do not invent unattested or unsupported combinations.
11. Output exactly one annotated utterance line and nothing else.

## B Training Label Inventory

Label	Label	Label	Label
[* m:++ed:i]	[* m:+ing]	[* m:base:er]	[* m:vsg:a]
[* m:++ed]	[* m:+s:a]	[* m:base:est]	[* m:vun:a]
[* m:++en:i]	[* m:+s]	[* m:base:s]	[* s:r:der]
[* m:++s:i]	[* m:θ's]	[* m:irr:ed]	[* s:r:gc:det]
[* m:++s]	[* m:θ3s:a]	[* m:irr:en]	[* s:r:gc:pro]
[* m:+3s:a]	[* m:θed]	[* m:irr:s]	[* s:r:prep]
[* m:+3s]	[* m:θing]	[* m:sub:ed]	
[* m:+ed:i]	[* m:θs:a]	[* m:sub:en]	
[* m:+ed]	[* m:=ed]	[* m:base:ed]	
[* m:+en]	[* m:=en]	[* m:base:en]	
[* m:=s]	[* m:allo]		

Table 6: Full training-set label inventory for the final confirmatory model package. The table lists all CHAT morphosyntactic error labels seen during training, irrespective of whether later evaluation support for a label comes from real or synthetic data.

## C Additional Evaluation Results

Split	N	EM	P	R	F1
Val.	370	82.4%	91.9%	88.7%	90.3%
Test	370	82.7%	94.5%	92.0%	93.2%

Table 7: Automatic evaluation on the synthetic support splits. Exact denotes full-line exact match against the gold annotated utterance; P, R, and F1 are micro-averaged over CHAT error tags.

Subset	N	Prev. exact	Reviewed	Both acc.	Model pref.	Post-hoc acc.
Test real (full)	687	643 (93.6%)	44	20	5	97.2%
Test real (labelled)	87	58 (66.7%)	29	11	3	82.8%

Table 8: Post-hoc adjudication results on test\_real. Previous exact gives the number of exact automatic matches before manual review. Post-hoc acceptable counts exact matches plus reviewed disagreement cases judged acceptable for both outputs or preferred for the model.

test_real					test_coverage				
Label	P	R	F1	N	Label	P	R	F1	N
[* m:allo]	100.0	100.0	100.0	14	[* m:=en]	100.0	100.0	100.0	12
[* m:vun:a]	100.0	100.0	100.0	2	[* m:irr:en]	100.0	100.0	100.0	11
[* s:r:der]	100.0	100.0	100.0	1	[* m:=ed]	100.0	100.0	100.0	10
[* m:=ed]	100.0	80.8	89.4	26	[* m:++en:i]	100.0	100.0	100.0	10
[* m:03s:a]	84.6	88.0	86.3	25	[* m:+en]	100.0	100.0	100.0	10
[* s:r:gc:pro]	87.5	70.0	77.8	10	[* m:0's]	100.0	100.0	100.0	10
[* m:base:ed]	75.0	75.0	75.0	4	[* m:base:en]	100.0	100.0	100.0	10
[* m:++ed:i]	100.0	50.0	66.7	2	[* m:base:er]	100.0	100.0	100.0	10
[* m:++ed]	50.0	100.0	66.7	1	[* m:base:est]	100.0	100.0	100.0	10
[* m:vsg:a]	50.0	50.0	50.0	2	[* m:irr:ed]	100.0	100.0	100.0	10
[* s:r:prep]	50.0	33.3	40.0	3	[* m:sub:ed]	100.0	100.0	100.0	10
[* m:+ed]	0.0	0.0	0.0	2	[* m:sub:en]	100.0	100.0	100.0	10
[* m:0ing]	0.0	0.0	0.0	2	[* m:++s]	90.9	100.0	95.2	10
[* s:r:gc:det]	0.0	0.0	0.0	2	[* m:+s:a]	90.9	100.0	95.2	10
[* m:0ed]	0.0	0.0	0.0	1	[* m:irr:s]	90.9	100.0	95.2	10
[* m:0s:a]	0.0	0.0	0.0	1	[* m:+ing]	90.0	90.0	90.0	10
					[* m:+3s]	81.2	100.0	89.7	13
					[* m:=s]	100.0	80.0	88.9	10
					[* m:base:s]	100.0	80.0	88.9	10
					[* m:++s:i]	88.9	80.0	84.2	10
					[* m:+3s:a]	100.0	60.0	75.0	10
					[* m:03s:a]	80.0	66.7	72.7	12

Table 9: Per-label automatic evaluation grouped by evaluation source. Values are percentages except for support (N). Real rows are taken from test\_real; synthetic rows are labels absent from test\_real and therefore reported on test\_coverage.

Label	Correct	N	Label	Correct	N
[* m:03s:a]	88.7%	328	[* m:sub:en]	100.0%	6
[* m:allo]	92.4%	105	[* s:r:prep]	83.3%	6
[* m:=ed]	78.6%	84	[* m:+ed]	100.0%	5
[* s:r:gc:pro]	71.4%	70	[* m:0s:a]	75.0%	4
[* m:vun:a]	68.2%	44	[* m:0's]	50.0%	4
[* m:vsg:a]	97.2%	36	[* m:irr:s]	50.0%	4
[* m:base:ed]	82.1%	28	[* m:=en]	33.3%	3
[* s:r:der]	23.8%	21	[* m:++s:i]	50.0%	2
[* m:irr:ed]	86.7%	15	[* m:+3s:a]	50.0%	2
[* m:++ed:i]	76.9%	13	[* m:++est]	100.0%	1
[* m:++ed]	61.5%	13	[* m:+en]	100.0%	1
[* m:0ing]	100.0%	11	[* m:base:en]	100.0%	1
[* m:+3s]	54.5%	11	[* m:irr:en]	100.0%	1
[* m:0ed]	80.0%	10	[* m:=ed:i]	0.0%	1
[* s:r:gc:det]	50.0%	10	[* m:+ing]	0.0%	1
[* m:sub:ed]	100.0%	9	[* m:base:der]	0.0%	1
[* m:+ing]	100.0%	8			

Table 10: Per-label human-reviewed exactness on the rest of the ENNI corpus, sorted by support. Because ENNI does not have exhaustive gold annotation, this table does not report recall or F1.

## D Expanded CHAT Scheme Reference

Level 1		Level 2		Level 3	
Code	Meaning	Code	Meaning	Code	Meaning
[* m:]	morphosyntactic error	0	missing	-ing	progressive
[* s:]	substitution error	base:	bare form	-3s	3SG
		sub:	substitution	-ed	past
		irr:	irregular	-en	perfective
		=	over-regularisation	-s	plural
		+	superfluous	's	possessive
		++	double marking	-s'	plural possessive
		vsg:	irregular verb 3SG	-er	comparative
		vun	irregular verb unmarked	-est	superlative
		allo	allomorph	a	agreement
		s:r:	related lexical substitution	i	irregular
		s:r:gc:	related grammatical substitution	POS	target POS

Table 11: Expanded reference table for CHAT-style error-label components used in this study.

# Author Index

- Accilien, Rose Flaure, 1  
Arikawa, Andrea, 59  
Bakuri-Jørgensen, Rohat, 187  
Basmov, Victoria, 257  
Bouamor, Houda, 146  
Chantova, Yasena, 113  
Chen, Run, 173  
Claridge, Claudia, 95  
Clary, John, 241  
Conhaim, Katherine, 178  
Datta, Mim, 173  
Elliott, Desmond, 187  
Er, Mustafa Erolcan, 75  
Flynn, Matthew, 229  
Friedrich, Annemarie, 95  
Gammeltoft-Hansen, Thomas, 187  
Gershuni, Hillel, 33  
Goldberg, Yoav, 257  
Hennhöfer, Oliver, 309  
Heyl, Marieke Anne, 187  
Hirschberg, Julia, 173  
Hladká, Barbora, 22  
Hossain, Md Sajjad, 241  
Humblot-Renaux, Galadrielle, 187  
Hussein, Abdelrahman Abdel Latif, 146  
Jahromi, Mohammad N. S., 187  
Jana, Abhik, 130  
Kahanda, Indika, 59  
Kinkel, Steffen, 309  
Kirilova, Tsvetina, 113  
Kraus, Ivan, 22  
Krückl, Xaver, 95  
LaManna, Ignazio, 173  
Lavadeep, Bandaru, 130  
Levine, Lauren, 178, 219  
Li, Lin, 241  
Lignos, Constantine, 229  
Litman, Diane, 158  
Liu, Yang Janet, 158  
Marcheva-Nash, Mila, 113  
Massoud, Walid, 146  
Meyer, Joel, 241  
Moeslund, Thomas B., 187  
Mompelat, Ludovic, 1  
Motalík Hodková, Kateřina, 22  
Murphy Høgenhaug, Anna, 187  
Newman, Sam, 229  
Obiso, Timothy, 229  
Pavlova, Ivelina, 113  
Perkins, Judy A., 241  
Prado, John, 173  
Puterbaugh, Ryan, 173  
Raghav, Ritwik, 130  
Schmück, Hanna, 95  
Shah, Vishwaa, 59  
Shin, Gyu-Ho, 12  
Shmidman, Avi, 33  
Stage Jarlner, Asta S., 187  
Stefanova, Tsvetelina, 113  
Strötgen, Jannik, 309  
Sun, Weiwei, 113  
Sung, Hakyung, 12  
Tao, Zihao, 173  
Tsarfaty, Reut, 257  
Tézil, David, 1  
Urban, Veronika, 95  
Vasileva, Yoana, 113  
Venturini, Shamira, 309  
Vlachou, Maria, 187  
Zaghouani, Wajdi, 47  
Zekri, Abdullah Mohamed Mohamed, 146  
Zeldes, Amir, 178, 219  
Zeman, Sonja, 95  
Zhong, Yang, 158