

# From Disagreement to Understanding: The Case for Ambiguity Detection in NLI

Chathuri Jayaweera and Bonnie J. Dorr

University of Florida, Gainesville, FL, USA

{chathuri.jayawee, bonniejdorr}@ufl.edu

## Abstract

This position paper argues that annotation disagreement in Natural Language Inference (NLI) is not mere noise but often reflects meaningful variation, especially when triggered by ambiguity in the premise or hypothesis. While underspecified guidelines and annotator behavior contribute to variation, content-based ambiguity provides a process-independent signal of divergent human perspectives. We call for a shift toward ambiguity-aware NLI that first identifies ambiguous input pairs, classifies their types, and only then proceeds to inference. To support this shift, we present a framework that incorporates ambiguity detection and classification prior to inference. We also introduce a unified taxonomy that synthesizes existing taxonomies, illustrates key subtypes with examples, and motivates targeted detection methods that better align models with human interpretation. Although current resources lack datasets explicitly annotated for ambiguity and subtypes, this gap presents an opportunity: by developing new annotated resources and exploring unsupervised approaches to ambiguity detection, we enable more robust, explainable, and human-aligned NLI systems.

## 1 Introduction

This paper takes a position on how disagreement in Natural Language Inference (NLI) is best understood and modeled. While prior work has often treated annotator disagreement as noise—something to be minimized or resolved (Snow et al., 2008; Bowman et al., 2015)—we argue that such disagreement can reflect meaningful, coexisting interpretations grounded in linguistic ambiguity.

NLI, also known as Recognizing Textual Entailment (RTE) (Dagan et al., 2005), aims to classify the relationship between a premise (P) and a hypothesis (H). Suppose  $P1=John\ likes\ Mary$ ,  $P2=John\ lives\ near\ Mary$ ,  $H1=John\ knows\ Mary$ ,  $H2=John\ doesn't\ know\ Mary$ . Standard inference




Figure 1: Framework for ambiguity-aware NLI: First detect whether a (P)remise or (H)ypothesis is ambiguous and, if so, classify the ambiguity type. Generate disambiguated versions and pass these to the inference classifier. Linguistic and other relevant background knowledge inform each stage. Gray = focus of this paper; white = supporting stages.

labels would assign *entailment* to (P1, H1), *contradiction* to (P1, H2), and *neutral* to (P2, H1). However, humans may diverge: (P1, H2) could also be neutral if *likes* is interpreted as distant admiration (e.g., of a celebrity).

We adopt a perspectivist reframing of NLI that treats variation in inference classifications not as a flaw but as an inherent feature of natural language understanding. We emphasize content-based ambiguity as a central source of disagreement and advocate for deeper exploration of its role in shaping inference judgments.

We situate our analysis within a framework for handling ambiguous NLI instances (Figure 1). This framework first determines whether an input pair is ambiguous; if so, it disambiguates the pair into dis-

tinct yet plausible human interpretations, enabling predictions aligned with each interpretation.

NLI is pivotal in understanding semantic relationships and is central to evaluating how well language models process natural language. NLI benchmarks are typically constructed using human-annotated entailment labels (Bowman et al., 2015; Williams et al., 2018). Despite frequent disagreements among annotators on the “correct” label for a given premise-hypothesis pair, most NLI research assumes a single “true” inference for each case.

Instances that deviate from this assumption are either filtered out during dataset construction (Bayer et al., 2005) or handled via majority vote (Bowman et al., 2015), based on the belief that annotation disagreements reflect random error rather than systematic variation. However, this approach contradicts the original purpose of NLI: to model what a reasonable, attentive, and informed human would plausibly infer from text (Manning, 2006).

Recent studies have challenged the assumption that annotation disagreements are mere noise, demonstrating instead that such disagreements exhibit reproducible patterns grounded in legitimate interpretive differences (Pavlick and Kwiatkowski, 2019). This recognition has motivated efforts to model the full distribution of plausible human inferences (Chen et al., 2020; Meissner et al., 2021).

While these efforts are important, understanding *why* such differences arise is equally critical for developing systems that reflect multiple human interpretations. This position paper argues for an NLI modeling goal that centers on **identifying and categorizing ambiguity into recurring interpretive patterns, rather than merely modeling annotator distributions to capture coexisting human perspectives**. We support this position through a review and analysis of existing research.

The next section reviews work on modeling annotator label distributions and their limitations, motivating a closer look at sources of disagreement in NLI. Section 3 examines prior categorizations of these sources, Section 4 highlights the unique role of ambiguity in premise-hypothesis pairs, and Section 5 surveys current ambiguity-focused NLI research and future directions.

## 2 Modeling annotator distribution

Most NLI benchmarks are constructed using human-annotated entailment labels, which often result in cases where multiple annotators assign dif-

ferent labels to the same premise-hypothesis pair. From the early days of NLI research, scholars have expressed concerns about how to handle such disagreements (Bayer et al., 2005).

Re-annotation of the RTE1 development and training sets reveals substantial discrepancies between the original and new labels (Bayer et al., 2005). Even after filtering out problematic examples, human judges only achieve a 91% agreement rate. Similar disagreements in the Stanford Natural Language Inference (SNLI) dataset complicate the process of learning robust decision boundaries for each entailment label (Pan et al., 2018).

Such cases are typically treated as “annotation noise,” resolved by assigning a majority label under the assumption that one “true” inference exists for each premise-hypothesis pair. However, growing evidence suggests that these disagreements reflect systematic, reproducible variation rather than random error (Pavlick and Kwiatkowski, 2019). In many cases, divergent annotations signal the existence of multiple plausible interpretations.

Current NLI models, trained on majority-labeled benchmarks, struggle to capture the full distribution of human judgments and tend to perform better when annotator agreement is high (Nie et al., 2020). This highlights both a dependence on agreement and a failure to model collective human reasoning. Meissner et al. (2021) further show that models trained on soft labels—distributions over annotator responses—better approximate human judgments and improve single-label prediction accuracy.

These findings have inspired a growing line of research focused on modeling human opinion distributions. For example, the Uncertain Natural Language Inference (UNLI) framework (Chen et al., 2020) proposes predicting subjective probabilities of entailment rather than coarse categorical labels. While UNLI captures a more probabilistic notion of inference, it targets average responses and does not attempt to model the full range of interpretations.

Zhang and de Marneffe (2021) contrast systematic inference (high agreement) and ambiguous cases (high disagreement). They build artificial annotators using BERT (Devlin et al., 2019) to simulate annotation variation, enabling downstream models to determine if a given premise-hypothesis pair is likely to elicit disagreement. Zhou et al. (2022) further improve modeling of opinion distributions beyond standard softmax assumptions.

Together, these studies mark a shift from a prescriptive view—assuming a single correct label—

toward a descriptive approach that acknowledges interpretive variations. They help pave the way for systems that capture ambiguity inherent in natural language. However, simply modeling disagreement alone does not explain *why* interpretations diverge. To advance beyond descriptive modeling, we argue that NLI systems must also systematically identify and categorize the sources of disagreement—especially content-based ambiguity—as a foundation for more perspective-sensitive inference (Plank, 2022). We next examine the sources that give rise to divergent judgments in NLI.

### 3 Disagreement sources in NLI

According to the “Triangle of Reference” (Aroyo and Welty, 2015), disagreement in annotation arises from three main sources: (1) interpretative ambiguity in the *input content* itself (Uncertainty in sentence meaning); (2) unclear annotation guidelines (Underspecification in guidelines); and (3) differences in annotators’ background knowledge or task understanding (Annotator behavior). This framework maps directly onto annotation workflows in NLI benchmarks. Building on this foundation, Jiang and de Marneffe (2022) propose a more fine-grained taxonomy for NLI, refining each category into subtypes that reflect recurring premise-hypothesis patterns (Figure 2).

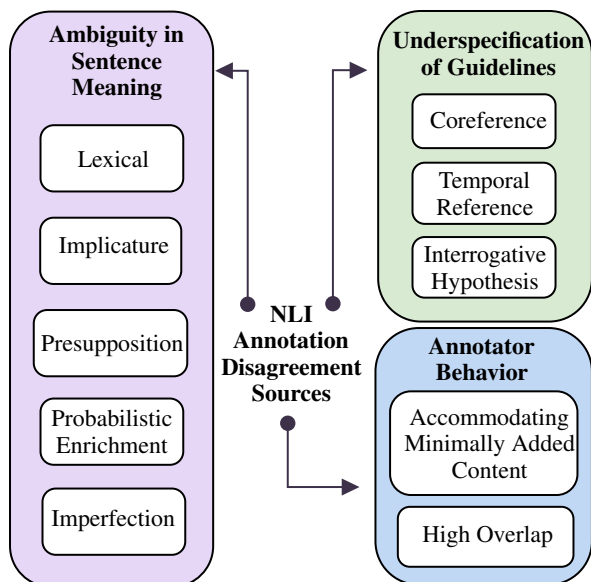


Figure 2: The taxonomy of disagreement sources developed by Jiang and de Marneffe (2022), building on Aroyo and Welty (2015). While these frameworks classify sources of disagreement, we argue for *reframing* such disagreement as a signal of coexisting interpretations to model—not noise to be resolved.

In the subsections below, we adopt a variant of this taxonomy, reframing *Uncertainty in Sentence Meaning* as *Ambiguity in Sentence Meaning* (Section 3.1), a shift already noted in Figure 2. This distinction is central to our position. For completeness, we also briefly describe the roles of guideline underspecification (Section 3.2) and annotator behavior (Section 3.3), though these are not the central emphasis of our position.

In addition, while we view Jiang and de Marneffe’s taxonomy as a valuable classification framework, we go further: rather than treating disagreement as noise to be explained or resolved, we reframe it as a meaningful signal of coexisting interpretations—something to be modeled directly as part of the NLI task.

#### 3.1 Ambiguity in Sentence Meaning

Ambiguity in sentence meaning—manifesting as multiple plausible interpretations—is a major source of disagreement in NLI annotations. In the taxonomy introduced by Jiang and de Marneffe (2022), this form of content-based ambiguity is further divided into five subtypes: Lexical, Implicature, Presupposition, Probabilistic Enrichment, and Imperfection. Together, these categories reflect the range of interpretative uncertainty that arises from the language content itself, independent of annotator knowledge or instructions specified in annotation guidelines.

Lexical arises when a word or phrase in the premise or hypothesis has multiple possible senses or is underspecified. Implicature refers to cases where the hypothesis expresses a logical or pragmatic implication of the premise, leaving room for divergent judgments depending on the reader’s perspective. Presupposition covers instances where the hypothesis draws on background presuppositions introduced by the premise, which may or may not be universally shared. Probabilistic Enrichment denotes cases where the inference relationship is not categorical, but depends on plausibility or likelihood, producing variation in individual perception (Figure 3). Imperfection includes typos, grammatical errors, or fragmented phrasing that impede clear interpretation.

While this categorization is based on a manually analyzed sample and shaped by the judgments of linguistically trained annotators, it offers a valuable foundation for surfacing and organizing patterns of interpretative variation in NLI. Although it does not capture the full range of ambiguities present in

#### Probabilistic Enrichment

**Premise:** “I think this report shows that we have had an inordinately productive and successful year.”  
**Hypothesis:** “The report shows that we need to be productive to have a successful year”

Figure 3: Probabilistic Enrichment ambiguity: Annotator’s label choice—Entailment or Neutral—depends on whether the relationship between productivity and success mentioned in the premise is considered plausible or not.

premise-hypothesis pairs, it provides an important starting point for tracing the roots of disagreement in NLI annotations and for recognizing how such divergences arise from legitimate differences in interpretation rather than annotation error.

### 3.2 Underspecification of Guidelines

Guideline underspecification is another source of annotation disagreement, but unlike content-based ambiguity, it often reflects task design flaws that can be addressed through clearer instructions. Even when the premise-hypothesis pairs are unambiguous, annotators may diverge in how they interpret or apply the labeling instructions if those instructions lack sufficient precision or fail to address edge cases. Jiang and de Marneffe (2022) identify three specific subtypes under this category: Coreference, Temporal Reference, and Interrogative Hypothesis.

#### Temporal Reference

**Premise:** “You wake up one bright autumn morning and you’re halfway to the subway when you decide to walk to work instead.”  
**Hypothesis:** “You wake up early and decide to walk instead of take the subway.”

Figure 4: Temporal Reference disagreement: Annotator’s label choice—Contradiction or Entailment—depends on whether the decision is interpreted as happening before or after the commute, respectively.

Coreference cases involve unclear assumptions about whether entities in the premise and hypothesis refer to the same thing. Without guidance on how strongly to assume shared reference, annotators may reach inconsistent labels.

Temporal Reference arises when it is unclear when the hypothesis should be evaluated. In Figure 4 one annotator might interpret the decision as occurring **before** the commute (Contradiction), while another may align it with a decision made **during** the commute (Entailment).

The Interrogative Hypothesis covers cases

where the hypothesis is phrased as a question. Since questions are not truth-apt (i.e., not directly true or false), annotators must infer an implied assertion. Jiang and de Marneffe (2022) focus only on interrogative hypotheses, while others (e.g., Gubelmann et al. (2023)) argue that interrogative premises can cause similar confusion.

These sources of disagreement are worth recognizing, but they stem from instruction gaps rather than genuine interpretive variation—and thus lie outside this paper’s primary focus.

### 3.3 Annotator Behavior

The third pillar of disagreement, according to the Triangle of Reference, is variation in annotation behavior: differences in background knowledge, beliefs, attention, or task interpretation across annotators. While often treated as noise in NLI pipelines, such variation can reflect meaningful differences in how people reason with language. Two annotators may bring different contextual assumptions to the same premise-hypothesis pair, leading to divergent but reasonable judgments.

Jiang and de Marneffe (2022) identify two specific behavioral tendencies that contribute to such disagreement: Accommodating Minimally Added Content and High Overlap. The first involves hypotheses that add a small amount of plausible but unstated information. Some annotators accept this as implied, while others reject it based on stricter entailment criteria. The second reflects a tendency to judge Entailment based on surface-level similarity—lexical or structural. This can lead some to overestimate entailment based on form rather than meaning, while others focus on more subtle semantic distinctions (Figure 5).

#### High Overlap

**Premise:** “The sunlight, piercing through the branches, turned the auburn of her hair to quivering gold.”  
**Hypothesis:** “The auburn of her hair became golden then the sunlight hit it.”

Figure 5: High Overlap disagreement: Annotators may infer Entailment between the premise and hypothesis due to the high lexical overlap, while the meaning of the two sentences suggests Contradiction.

Some annotation tendencies stem not from errors, but from genuine interpretive variation. For example, Accommodating Minimally Added Content reflects meaningful differences, whereas High Overlap more likely signals annotation error. This underscores the importance of evaluating

annotator behavior carefully, rather than assuming that all variation reflects valid perspectives.

Among the various sources, content-based ambiguity is the most direct and reliable indicator of genuine interpretive divergence. While understanding annotator behavior is useful, our focus is on detecting ambiguity in the language itself. Even so, recognizing behavioral patterns can inform future perspectivist NLI systems that accommodate multiple interpretations. Next, we motivate why content-based ambiguity merits special attention relative to guideline and annotator effects.

#### 4 Why Does Content-Based Ambiguity Deserve Special Attention?

Among the disagreement sources outlined above, content-based ambiguity stands out as the only type that can be systematically addressed through computational modeling without relying on additional information about the annotators or the guidelines they follow. This form of ambiguity originates from the text itself, independent of the annotation process, yet it remains a fundamental driver of divergent interpretations. As such, it represents a root cause of disagreement inherent to natural language, posing a persistent challenge for inference systems aiming for consistent and reliable predictions.

Implicature Ambiguity
<b>Premise:</b> “It hopes to bring on another 25 or 35 people when the new building opens next fall.”
<b>Hypothesis:</b> “They already have a waiting list for the new building”

Figure 6: Implicature Ambiguity: Annotators may infer a waiting list from *hopes to bring on another 25 or 35 people*. If so, they label it Entailment; if not, they label Neutral.

Consider the Implicature ambiguity in Figure 6. Some annotators interpret *hopes to bring on another 25 or 35 people* as implying a *waiting list* and choose Entailment. Others focus strictly on what is stated and select Neutral. This interpretative variability stems from linguistic ambiguity rather than annotator background or faulty instructions. Such cases underscore the importance of treating content-based ambiguity as central to analysis, rather than dismissed as noise.

Jiang and de Marneffe (2022)’s findings indicate that the most common sources of disagreement fall under content-based ambiguity, underscoring its prevalence. In contrast, issues related to underspecified guidelines can typically be resolved through

clearer instructions and better annotation practices, meaning they do not strongly reflect genuine differences in human interpretation.

Similarly, while some annotator behaviors—such as Accommodating Minimally Added Content—reflect natural variation, others like High Overlap may undermine the goals of the NLI task. Disagreements stemming from annotator behavior or guideline underspecification therefore warrant scrutiny before being treated as meaningful. Not all disagreements are noise, though some clearly reflect error (Weber-Genzel et al., 2024). In contrast, content-based ambiguity arises from the language itself and requires no external filtering or supervision, making it a uniquely reliable source of interpretive variation. Identifying such ambiguity supports the creation of disambiguated versions, allowing NLI benchmarks to better capture the range of plausible interpretations. Beyond benchmarking, ambiguity-aware modeling has practical consequences in downstream settings.

Identifying content-based ambiguity in NLI data has significant real-world implications. NLI frequently serves as a core component of fact-verification pipelines, where it is used to assess the relationship between claims and supporting evidence (Thorne et al., 2018; Jayaweera et al., 2024). Effectively pinpointing potential ambiguities—including those intentionally introduced—strengthens such pipelines by improving their capacity to flag potentially misleading content in real-world settings (Liu et al., 2023).

However, there are currently no established methods for disentangling disagreements caused by content-based ambiguity from those arising due to underspecified annotation guidelines or annotator behavior. As a result, most existing work focuses primarily on detecting premise-hypothesis pairs with high annotation disagreement, rather than investigating the underlying types of disagreement, particularly those stemming from ambiguity. Therefore, there is a necessity to build models that: (1) identify ambiguous premise-hypothesis pairs and (2) classify the respective ambiguity type. These observations motivate two concrete tasks—ambiguity detection and ambiguity classification—which we discuss next.

#### 5 Understanding Ambiguity in NLI

NLI systems aim to determine the inference relationship between a given premise and hypoth-

esis, but ambiguity in either can complicate that process (Figure 6). This often leads to discrepancies among annotators, who may assign different inference labels based on their individual interpretations. In some cases, annotators may even agree on the same label while interpreting the text differently—a phenomenon known as within-label variation (Jiang et al., 2023). Further complicating matters, ambiguity may arise in the premise, the hypothesis, or both, increasing the complexity of inference decisions (Liu et al., 2023).

While some efforts have been made to develop models that detect instances with high annotator disagreement (Jiang and de Marneffe, 2022; Jiang et al., 2023; Park and Kim, 2025), there are no existing implementations that specifically identify or classify ambiguous instances in NLI—underscoring the need for systems designed to address this gap.

### 5.1 Ambiguity Detection in NLI

We define *ambiguity detection* in NLI as identifying instances that elicit divergent interpretations due to input ambiguity—whether in the premise, the hypothesis, or both.

Jiang and de Marneffe (2022) explore the detection of high-disagreement instances in NLI using multi-label prediction and a four-class classification scheme (Entailment, Contradiction, Neutral, and Complicated). However, their work does not go further to distinguish ambiguity as a specific cause of disagreement. Jiang et al. (2023) build on this by incorporating explanations for disagreement but still focus solely on identifying highly contested instances.

Park and Kim (2025) attempt to detect ambiguous cases in NLI benchmarks using hidden layer representations of Large Language Models (LLMs), but their training data includes disagreements from all categories, making the system a general disagreement detector rather than a model focused on ambiguity. Liu et al. (2023) assess language models’ ability to detect ambiguous instances using the Ambient dataset, but their results show that model performance remains below human-level accuracy.

These studies reflect the current state of ambiguity detection in NLI, highlighting the need for further investigation. A key challenge in developing systems to identify ambiguous premise-hypothesis pairs is the lack of datasets annotated for ambiguity. Creating annotated datasets and exploring unsupervised methods are essential next steps.

To address the current scarcity of ambiguity-type annotated NLI data, we leverage existing datasets that already incorporate disambiguations (Liu et al., 2023) and explanations (Jiang et al., 2023) as assistive cues to annotate ambiguity types. This approach would help create a more cohesive dataset that integrates insights across the various taxonomies discussed in Section 5.

At the same time, the limited scale of these resources highlights the need for additional strategies. Promising directions include data augmentation techniques such as paraphrasing, the continued use of manual annotation to ensure high-quality gold standards, and the strategic use of large language models (LLMs) as evaluators. Together, these methods can substantially expand the availability of annotated data, enabling both broader coverage of ambiguity types and more robust evaluation of ambiguity-aware NLI systems.

### 5.2 Ambiguity Classification in NLI

*Ambiguity classification* identifies the exact type(s) of ambiguity present in a premise-hypothesis pair. Several taxonomies have been developed to categorize the various forms of ambiguity found in NLI inputs. As noted and illustrated in Figure 2, Jiang and de Marneffe (2022) present a taxonomy comprising five ambiguity types. These have been identified in samples from the ChaosNLI (Nie et al., 2020) and MNLi (Williams et al., 2018) datasets.

Liu et al. (2023) introduce a taxonomy based on expert linguistic annotations of the Ambient dataset, which contains both curated and generated ambiguous premise-hypothesis pairs. They identify additional ambiguity types in NLI data, including Syntactic, Pragmatic, Scopal, and Figurative ambiguities, while grouping others under a residual Other category.

Building on this, Li et al. (2024) refine the classification by proposing finer-grained types such as Type/Token and Collective/Distributive, and aligning with Jiang and de Marneffe (2022) through the inclusion of Presupposition and Implicature. These refinements reveal further unexplored ambiguities that enhance the understanding of human interpretations. Drawing on these developments, we present a unified taxonomy,<sup>1</sup> that organizes ambiguity types into four broad categories—Lexical, Syntactic, Semantic, and Pragmatic—to support a more comprehensive view (Figure 7).

<sup>1</sup>Refer to (Li et al., 2024) for the definitions of each ambiguity type not described in this paper.

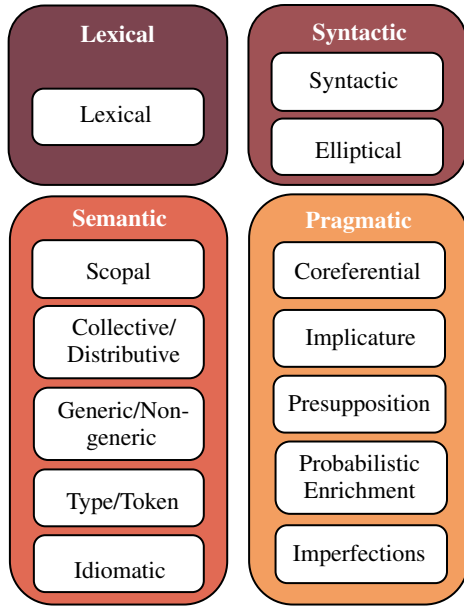


Figure 7: Unified ambiguity type taxonomy: We build on prior taxonomies (Jiang and de Marneffe, 2022; Liu et al., 2023; Li et al., 2024), organizing them into four broad types—Lexical, Syntactic, Semantic and Pragmatic—to support tailored detection strategies based on common characteristics.

However, to our knowledge, no existing system automatically identifies the ambiguity type(s) present in a given premise–hypothesis pair. This gap highlights the practical importance of our framework: by organizing ambiguity types into linguistically grounded categories, we lay the foundation for developing detection methods tailored to each type. In doing so, we advance toward more nuanced, interpretable NLI models that not only detect ambiguous input, but also explain *how* and *why* human interpretations diverge.

## 6 Call to Action: Recognizing Ambiguity as Signal, Not Noise

Disagreement among annotators in Natural Language Inference (NLI) is often treated as noise—something to minimize or discard. However, many of these disagreements reflect genuine interpretive differences, often triggered by ambiguity in the premise, hypothesis, or both. Our analysis suggests that while underspecified annotation guidelines and inconsistent annotator behavior can lead to label disagreement, such cases must be carefully scrutinized to distinguish between annotation errors and genuine differences in human interpretation.

In contrast, ambiguity in the NLI input itself—whether lexical, syntactic, semantic, or pragmatic—serves as a clear, process-independent signal of

interpretive variation, providing a basis for understanding how meaning can diverge across readers. This highlights a needed shift: from optimizing for annotator consensus to explicitly identifying and characterizing ambiguity as a central feature of natural language.

To support this shift, we outline the following two key directions:

- **Identify Ambiguous Pairs:** Develop robust methods to detect premise–hypothesis pairs that exhibit inherent ambiguity, using cues from linguistic theory, annotation patterns, and interpretability tools.
- **Classify Ambiguity Types:** Design strategies for distinguishing among different types of ambiguity. A unified classification framework that groups ambiguity types based on their shared characteristics can offer a foundation for designing targeted identification methods.

A key novelty of this work lies in articulating a unified framework (Figure 1) that extends beyond prior approaches focused solely on modeling annotation distribution. Whereas earlier efforts largely stop at detecting high-disagreement instances, this framework explicitly distinguishes ambiguity from other sources of variation by leveraging linguistic features and pertinent background knowledge.

The framework consists of four stages. The first determines whether a premise-hypothesis pair is inherently ambiguous, thereby distinguishing between genuine interpretive variation from annotation noise and other sources of disagreement. If an instance is ambiguous, the second classifies the type(s), creating systematic linkages to linguistic background knowledge.

The third stage generates relevant disambiguated versions, after which the fourth models inference classification for both ambiguous and non-ambiguous instances, based on predictions from earlier stages. The framework’s strength lies in offering a structured and operational foundation for ambiguity-aware NLI, moving the field from descriptive accounts of disagreement toward a principled methodology that can be empirically tested.

By pursuing these goals, we can build NLI models that are not only more aligned with human interpretation, but also more explainable in predictions. Ambiguity-aware systems better align with human interpretation and produce more consistent, interpretable, and robust predictions.

This reframing is not only timely—it is essential for developing NLI systems that reflect the complexity of human understanding, rather than abstracting it away.

While we advocate for a shift toward ambiguity-aware NLI systems, realizing this vision is currently constrained by a key limitation: the lack of datasets that are explicitly annotated for ambiguity and categorized by ambiguity type. Most existing NLI datasets are not designed with interpretive variation or ambiguity classification in mind, making it difficult to systematically identify and analyze ambiguous instances or to evaluate models on their ability to handle them.

This gap limits the development and benchmarking of methods for detecting and classifying ambiguity. The absence of gold-standard annotations for different ambiguity types hinders progress in training and evaluating models that aim to align more closely with human interpretive processes.

To address this, we suggest two complementary directions. First, there is a clear need for the **creation of new datasets** specifically annotated for ambiguity presence and type. Such resources would lay the groundwork for both empirical analysis and model development. Second, we see promise in **exploring unsupervised or weakly supervised methods** that can surface potential ambiguities without requiring extensive manual labeling. Techniques leveraging patterns of annotator disagreement, discourse features, or model uncertainty could offer scalable alternatives in the absence of annotated data.

Despite current limitations, these strategies offer a promising path toward building NLI systems that better reflect the complexity and nuance of human language understanding.

## Limitations

The framework we articulate for ambiguity-aware NLI systems establishes a theoretical foundation, with empirical validation remaining an important next step. As a position paper, our aim is to stimulate discussion and motivate future empirical work. The framework’s logical rigor and integration of existing taxonomies offer a strong basis for future experimentation and evaluation.

Future research must address the creation of datasets explicitly annotated for ambiguity, alongside the development and evaluation of systems to identify ambiguous instances in NLI. Such efforts

will contribute to a deeper understanding by identifying indicators of different ambiguity types, and characterizing how they shape inference judgments. We also anticipate exploring hybrid approaches that combine linguistic analysis with large language models to advance ambiguity detection for NLI.

## Acknowledgments

This work would not have been possible without the generous startup support provided by Dr. Herbert Wertheim through the Herbert Wertheim College of Engineering at the University of Florida.

## References

- Lora Aroyo and Chris Welty. 2015. [Truth Is a Lie: Crowd Truth and the Seven Myths of Human Annotation](#). *AI Magazine*, 36(1):15–24. Number: 1.
- Samuel Bayer, John Burger, Lisa Ferro, John Henderson, and Alexander Yeh. 2005. MITRE’s Submissions to the EU Pascal RTE Challenge. In *Proceedings of the Pattern Analysis, Statistical Modelling, and Computational Learning (PASCAL) Challenges Workshop on Recognising Textual Entailment*, pages 41–44.
- Samuel R. Bowman, Gabor Angeli, Christopher Potts, and Christopher D. Manning. 2015. [A large annotated corpus for learning natural language inference](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 632–642, Lisbon, Portugal. Association for Computational Linguistics.
- Tongfei Chen, Zhengping Jiang, Adam Poliak, Keisuke Sakaguchi, and Benjamin Van Durme. 2020. [Uncertain Natural Language Inference](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8772–8779, Online. Association for Computational Linguistics.
- Ido Dagan, Oren Glickman, and Bernardo Magnini. 2005. [The PASCAL recognising textual entailment challenge](#). In *Proceedings of the First international conference on Machine Learning Challenges: evaluating Predictive Uncertainty Visual Object Classification, and Recognizing Textual Entailment, MLCW’05*, pages 177–190, Berlin, Heidelberg. Springer-Verlag.
- 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.
- Reto Gubelmann, Aikaterini-lida Kalouli, Christina Niklaus, and Siegfried Handschuh. 2023. [When Truth Matters - Addressing Pragmatic Categories in](#)



- Natural Language Inference (NLI) by Large Language Models (LLMs). In *Proceedings of the 12th Joint Conference on Lexical and Computational Semantics (\*SEM 2023)*, pages 24–39, Toronto, Canada. Association for Computational Linguistics.
- Chathuri Jayaweera, Sangpil Youm, and Bonnie J Dorr. 2024. **AMREx: AMR for Explainable Fact Verification**. In *Proceedings of the Seventh Fact Extraction and VERification Workshop (FEVER)*, pages 234–244, Miami, Florida, USA. Association for Computational Linguistics.
- Nan-Jiang Jiang and Marie-Catherine de Marneffe. 2022. **Investigating Reasons for Disagreement in Natural Language Inference**. *Transactions of the Association for Computational Linguistics*, 10:1357–1374. Place: Cambridge, MA Publisher: MIT Press.
- 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.
- Margaret Y. Li, Alisa Liu, Zhaofeng Wu, and Noah A. Smith. 2024. **A Taxonomy of Ambiguity Types for NLP**. *arXiv preprint*. ArXiv:2403.14072 [cs].
- Alisa Liu, Zhaofeng Wu, Julian Michael, Alane Suhr, Peter West, Alexander Koller, Swabha Swayamdipta, Noah Smith, and Yejin Choi. 2023. **We’re Afraid Language Models Aren’t Modeling Ambiguity**. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 790–807, Singapore. Association for Computational Linguistics.
- Christopher D. Manning. 2006. **LOCAL TEXTUAL INFERENCE : IT’S HARD TO CIRCUMSCRIBE , BUT YOU KNOW IT WHEN YOU SEE IT - AND NLP NEEDS IT**.
- Johannes Mario Meissner, Napat Thumwanit, Saku Sugawara, and Akiko Aizawa. 2021. **Embracing Ambiguity: Shifting the Training Target of NLI Models**. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 862–869, Online. 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.
- Boyuan Pan, Yazheng Yang, Zhou Zhao, Yueting Zhuang, Deng Cai, and Xiaofei He. 2018. **Discourse Marker Augmented Network with Reinforcement Learning for Natural Language Inference**. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 989–999, Melbourne, Australia. Association for Computational Linguistics.
- Hancheol Park and Geonmin Kim. 2025. **Where do LLMs Encode the Knowledge to Assess the Ambiguity?** In *Proceedings of the 31st International Conference on Computational Linguistics: Industry Track*, pages 445–452, 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. Place: Cambridge, MA Publisher: MIT Press.
- 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.
- Rion Snow, Brendan O’Connor, Daniel Jurafsky, and Andrew Ng. 2008. **Cheap and Fast – But is it Good? Evaluating Non-Expert Annotations for Natural Language Tasks**. In *Proceedings of the 2008 Conference on Empirical Methods in Natural Language Processing*, pages 254–263, Honolulu, Hawaii. Association for Computational Linguistics.
- James Thorne, Andreas Vlachos, Christos Christodoulopoulos, and Arpit Mittal. 2018. **FEVER: a large-scale dataset for Fact Extraction and VERification**. *arXiv preprint*. ArXiv:1803.05355.
- 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.
- Adina Williams, Nikita Nangia, and Samuel Bowman. 2018. **A Broad-Coverage Challenge Corpus for Sentence Understanding through Inference**. 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 1112–1122, New Orleans, Louisiana. Association for Computational Linguistics.
- Xinliang Frederick Zhang and Marie-Catherine de Marneffe. 2021. **Identifying inherent disagreement in natural language inference**. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4908–4915, Online. Association for Computational Linguistics.

Xiang Zhou, Yixin Nie, and Mohit Bansal. 2022. **Distributed NLI: Learning to Predict Human Opinion Distributions for Language Reasoning**. In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 972–987, Dublin, Ireland. Association for Computational Linguistics.