

# DeDisCo at the DISRPT 2025 Shared Task: A System for Discourse Relation Classification

Zhuoxuan Ju, Jingni Wu, Abhishek Purushothama, Amir Zeldes

Corpling Lab

Georgetown University

{zj153, jw2175, ap2089, amir.zeldes}@georgetown.edu

## Abstract

This paper presents **DeDisCo**, Georgetown University’s entry in the DISRPT 2025 shared task on discourse relation classification. We test two approaches, using an mt5-based encoder and a decoder based approach using the openly available Qwen model. We also experiment on training with augmented dataset for low-resource languages using matched data translated automatically from English, as well as using some additional linguistic features inspired by entries in previous editions of the Shared Task. Our system achieves a macro-accuracy score of 71.28, and we provide some interpretation and error analysis for our results.<sup>1</sup>

## 1 Introduction

Recent computational work on discourse relations has introduced richer multilingual corpora (Liu et al., 2024; Zeldes et al., 2025) and advanced transformer-based methods for implicit and explicit discourse relation classification (e.g. Li et al. 2024; Metheniti et al. 2024). While most previous studies focus on implicit relations within a single language (Liu and Strube, 2023), the DISRPT shared-task setting requires handling both explicit and implicit relations across multiple languages and annotation frameworks (Braud et al., 2024). This means that for example explicit CAUSAL relations as in 1, where the relation can be identified thanks to the explicit connective ‘because’, are targeted next to implicit ones, as in 2, where it is implied that it’s hard for the speaker to think about food (because) they are not hungry (both example taken from eng.erst.gum).

(1) [I bought it]<sub>unit1</sub> [**because** it was funny]<sub>unit2</sub>

<sup>1</sup>Full disclosure: our team includes both organizers and dataset annotators from the shared task. Results were reproduced independently by other organizers. All code is available at <https://github.com/gucorpling/disrpt25-task>.

(2) [I’m so not hungry right now,]<sub>unit1</sub> [it’s hard for me to think about food.]<sub>unit2</sub>

In DISRPT 2025 (see Braud et al. 2025), we designed and tested two architectures – one based on an encoder model and one based on a decoder – and submitted the decoder-only system, which is called DeDisCo (**D**ecoder-based **D**iscourse **C**ognoscente), which was trained via supervised fine-tuning with instruction-style prompts, tailored for multilingual discourse relation classification. Given the shared task constraint of a single multilingual model, supervised fine-tuning allowed us to efficiently fine-tune a compact decoder architecture while maintaining robust generalization across languages. Although we also explored encoder-based alternatives for methodology comparison, the decoder model formed the core of our official submission.

## 2 Related Work

### 2.1 Discourse Frameworks and Datasets

Multiple frameworks have been used to analyze and organize natural language discourse such as RST (Mann and Thompson, 1988) and eRST (Zeldes et al., 2025), PDTB (Marcu and Wong, 2004), SDRT (Lascarides and Asher, 2007), Discourse Dependencies (Stede et al., 2016), and ISO 24617-8 (Tomaszewska et al., 2024). These frameworks apply different structural and semantic guidelines for this purpose. However all of these frameworks utilize components such as discourse units and relations to build either shallow, or hierarchical/graph based discourse structures. Discourse units are then combined with *discourse relations labels* to describe the complex discourse structures present in texts.

Considerable efforts have been dedicated to building discourse corpora in multiple languages containing various text types. Combined with the variety of frameworks, many corpora contain fea-

tures specific to languages, genres, and frameworks: for example, in RST, most datasets use clauses as discourse units, but only some datasets allow for discontinuous units; in PDTB, some datasets focusing on implicit relations only allow relations between sentences, while in others, smaller units are included. Additionally, the nature of the text will make some relations more or less prevalent, such as questions (the QUERY relation), which appear almost only in dialogic text types. Hence, both language and genre-specific considerations can be applied in many instances, leading to heterogeneity in the data parameterized by the language, framework, corpus, and genre. Previous DISRPT Shared tasks (Zeldes et al., 2019, 2021; Braud et al., 2023) have brought together collections of these varied datasets, with the previous iteration (DISRPT 2023 Braud et al., 2023) containing 26 corpora annotated across 4 frameworks and 13 languages, and the 2025 edition bringing these up to 39 datasets in coming from 6 frameworks and 16 languages. New in this edition is also the unification of the label set to the same 17 labels across all datasets; however, labels may exhibit subtle differences in usage across corpora, meaning once again that encoding dataset specific input may be crucial for high performance across benchmarks.

## 2.2 Discourse Parsing and Relation Classification

The task of discourse parsing refers to constructing discourse structures (in a given scheme) from natural language text. This involves sub-tasks such as segmenting discourse units, connecting them in order, identifying lexical units that signal relations (in PDTB and eRST), and assigning relation labels. In addition to the inherent challenges of these tasks, heterogeneity leads to systems that are useful for only a small subset of the corpora. With neural models, this challenge is more severe due to the need for substantial amounts of training data to build such systems. The multilingual and multi-track nature of the DISRPT shared task may alleviate some of these problems by posing the problem as a multilingual one across all datasets, allowing low-resource languages to gain in performance from the additional data available in high-resource languages. The tracks on discourse unit segmentation, connective identification, and relation classification also allow systems to focus on these problems in isolation – in this paper we target only the latter task of relation classification, given

the units connected by the relation.

Regardless of the labels used, which are commonly framework-specific (Hovy, 1990), relation classification can be considered a single-label classification task over two ordered, non-overlapping textual input units. Multiple methods have been applied to the task, including in the context of full hierarchical discourse parsing, such as shift-reduce parsing (Marcu, 1999), feature based and neural parsers (Ji and Eisenstein, 2014), and simple span-based transformer encoder (Gessler et al., 2021; Metheniti et al., 2024), and decoder (Anuranjana, 2023) models.

## 2.3 Feature Encoding and Data Augmentation

Several past top systems Yu et al. (2019); Gessler et al. (2021); Metheniti et al. (2023) used hand-crafted features to mitigate discourse parsing challenges, even with transformer-based models. These range from corpus-level features such as language and framework, document-level features such as length or genre, and sample- or unit-level features such as lexical overlap or position. Unique features like whether two units share the same speaker were also crafted in corpus-specific systems. Features can be represented categorically (e.g., one-hot), as dense embeddings in neural architectures, or as plain text in LLM prompts.

Data augmentation has been used to improve both dataset and language-specific performance. Task-specific augmentation entails transforming or synthesizing data that is similar to the target dataset based on a distribution that also adheres to the dataset design. e.g. Liu et al. (2023) grouped together data for multiple languages with smaller corpora for such augmentation. For language-specific augmentation, the same or similar source for different languages can be translated into the target language, providing higher training data for the target language. This is popularly known as the translate-train paradigm (Conneau et al., 2018), which we employ below.

## 3 Data and Approach

The Discourse Relation Classification task of the DISRPT 2025 aims to classify discourse relations across a diverse set of languages and annotation frameworks. The relation classification task includes 38 of the joint task’s 39 corpora, spanning 16 languages and 6 different frameworks, with a unified set of 17 labels provided. For convenience, the language, corpora, and framework are part of

the listed in appendix D (Table 6). The majority of the data are annotated using Rhetorical Structure Theory (RST; Mann and Thompson, 1988) and the Penn Discourse Treebank framework (PDTB; Marcu and Wong, 2004). A smaller subset of the corpora is annotated using SDRT, discourse dependencies (DEP), eRST, and the ISO frameworks.

For our system, we experimented with both encoder-based and decoder-based models. After evaluating their performance, we selected the stronger decoder-based model for our official submission. Below, we described how each model was implemented with our feature set and data augmentation.

### 3.1 Features

**Language Corpus Framework (LCF)** As noted, the training data spans multiple annotation frameworks, corpora, and languages, which poses a challenge for generalization since relations defined under different schemes or languages are not always directly comparable. To address this diversity, we incorporate metadata into the input to help the model distinguish and generalize across datasets. We refer to these elements collectively as LCF features (Language, Corpus, Framework). For example, the dataset identifier `zho.rst.gcdt` indicates a Chinese corpus (`zho`) named GCDT (Peng et al., 2022b), annotated using the RST framework.

**DiscoDisco Features** The DiscoDisco system (Gessler et al., 2021) introduced several hand-crafted discourse features extracted from the data for the relation classification task (see Table 1), which were later shown to be highly effective (Metheniti et al., 2023). We selectively incorporated a subset of these features into our models, see Table 5 in appendix B for a detailed list of which DiscoDisco features are used in the decoder and the encoder separately. For exact details of these features we refer to the original DiscoDisco paper and system implementation.

**Direction** Discourse relations are marked between two segments of text, referred to as Arg1 and Arg2. In the DISRPT datasets, these segment pairs follow the text’s original sequence, and an extra column specifies the intended argument direction for the annotated relation (e.g.,  $1>2$ ). For example in a cause relation, the cause points to the result, but may appear first ( $1>2$ ) or second ( $1<2$ ). This directional information is incorporated into both of our models, but encoded in different ways (see below).

**Context** Text context beyond the two units being classified plays a crucial role in discourse relation classification. Notably, Dai and Huang (2018) demonstrated that paragraph-level context significantly enhances the prediction of implicit discourse relations. In our models, we explore adding context to the model input and the impact of varying context window sizes surrounding the target sentence.

To clarify the experimental setup, Table 4 in appendix B provides an overview of which features, context, and data augmentation used in each model configuration.

### 3.2 Data Augmentation

Given the limited size of training data for certain languages, many of which also exhibit lower model accuracy, we apply targeted data augmentation to enhance performance. We focus on six low-resource languages: Czech (14.6K tokens), Dutch (24.9K), French (32.7K), Basque (45.7K), German (66K), and Persian (67K), covering a total of seven datasets. While these six corpora are not strictly the smallest by token count, they exemplify low-resource conditions due to the combination of restricted training data for the entire language, and comparatively weak baseline performance. We therefore targeted them for augmentation to mitigate these weaknesses. Our augmentation strategy involves translating English training instances from a source corpus into the target languages using API calls to the ChatGPT 4.1 model (OpenAI et al., 2024), and providing this data for system training and replication following the shared task rules (our system does not access ChatGPT in any way at training or testing time).

To ensure compatibility with the target language, we select English data for translation based on four criteria: annotation framework, discourse relation label distribution, genre alignment, and overall dataset size. The detailed correspondences are provided in Table 2. For each target corpus, we generated augmented data equivalent to approximately 75% of its original size. This ratio was chosen to enrich the training set without overshadowing the signals from the original in-domain examples. To ensure the quality and relevance of the synthetic data, we implemented a multi-faceted filtering strategy. First, we maintained a label distribution in the augmented set that closely mirrored the original. Second, we aligned the data’s genre; for instance, as the German `deu.rst.pcc` contains mostly editorial texts and news, we primarily drew

Feature	Type	Ex.	Description
Genre	Cat.	reddit	Document genre (e.g., eng.erst.gum)
Children	Num.	2	Child units each discourse unit has
Discontinuous	Cat.	false	Unit tokens not contiguous
Is Sentence	Cat.	true	Unit is a complete sentence
Length Ratio	Num.	0.3	Token length ratio (u1 vs. u2)
Same Speaker	Cat.	true	Same speaker for u1 and u2
Doc. Length	Num.	214	Document length in tokens
Position	Num.	0.4	Unit position in doc (0–1)
Distance	Num.	7	Other units between u1 and u2
Lexical Overlap	Num.	3	Shared non-stoplist words

**Table 1:** Features used in 2021 DiscoDisco system.

Target Corpus	Source Corpus	Selected Source Genres
ces.rst.crdt	eng.erst.gum	essay, news
deu.pdtb.pcc	eng.pdtb.gum	essay, news, speech
deu.rst.pcc	eng.erst.gum	essay, news, speech
eus.rst.ert	eng.erst.gum	textbook, academic
fra.rst.prstc	eng.erst.gum	news, academic
nld.rst.nldt	eng.rst.(oll, sts)	bio, news, letter
fas.rst.prstc	eng.rst.rsttd	all

**Table 2:** Source–Target Genre & Framework Alignment for Translation-Based Data Augmentation

source material from the ‘essay’ and ‘news’ genres from the English GUM corpus, supplementing it with a small amount of ‘speech’ data to reach the target volume. Finally, we tried to enforce annotation guideline consistency. For example, we observed that German RST annotations do not segment relative clauses, unlike its counterpart in English. Therefore, we excluded any source examples with these incompatible structural patterns from the German augmentation set.

### 3.3 Pruned Qwen3-4B Decoder Only

For our decoder-only approach, we frame discourse relation classification as a generative task. Specifically, we feed a prompt to a decoder-only model, instructing it to directly select the correct label from a predefined label set included within the prompt itself. We employ the Qwen3-4B model<sup>2</sup> (Yang et al., 2025), chosen for its strong multilingual capabilities, supporting over 100 languages and dialects, which aligns well with the multilingual classification task. We apply supervised fine-tuning with instruction-style prompts to improve its task-specific performance.

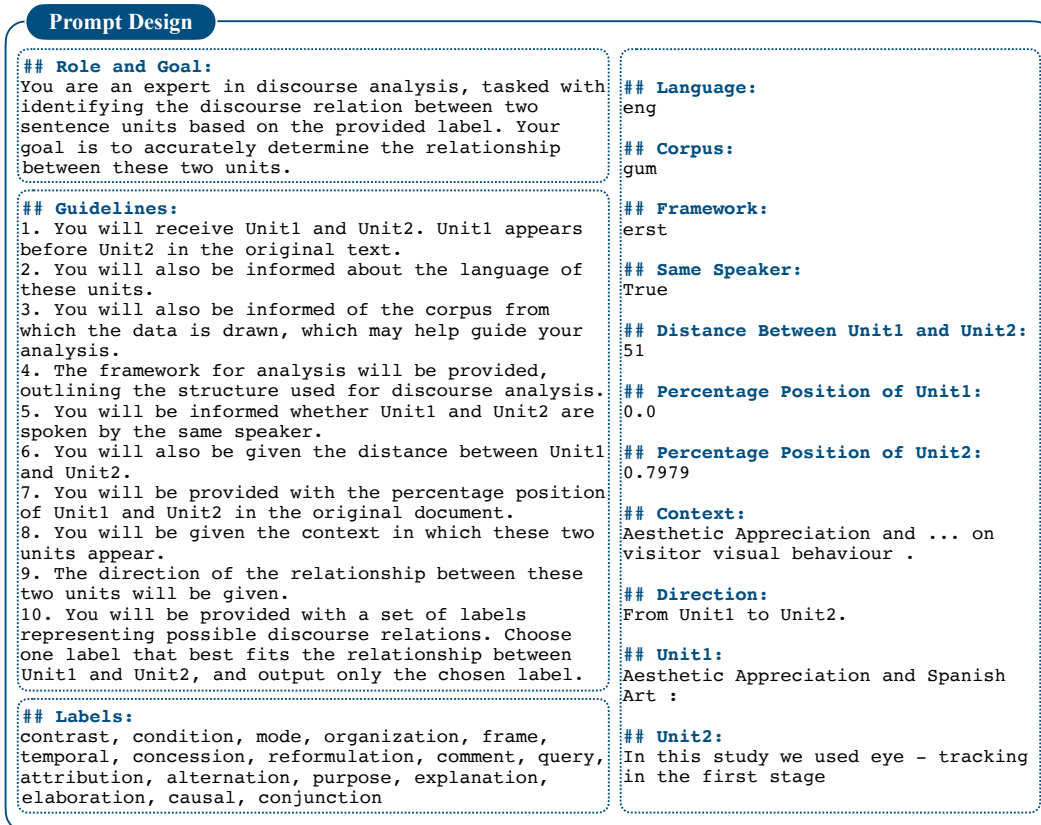
**Pruning** The public Qwen3-4B model originally contains 4.02 billion parameters, slightly exceeding the shared task’s 4B parameter limit. To address this, we adopt a pruning strategy based on layer

removal as proposed by Men et al. (2024), which identifies redundant layers by measuring the similarity between their input and output representations. We determined that removing a single, most redundant layer was sufficient to meet the parameter requirement. After fine-tuning, the resulting pruned model<sup>3</sup> achieved performance on par with its unpruned counterpart.

**Supervised Fine-Tuning** Our methodology involves full-parameter supervised fine-tuning of the model on the task-specific dataset, which is reformulated into an instruction-style prompt format. Each instruction is enriched with a comprehensive set of features, including LCFs, direction, context, and selected DiscoDisco features (e.g., same speaker, position), as detailed in Section 3.1. The context is constructed from the sentence immediately preceding the first argument, the sentence(s) containing both arguments, and the sentence immediately following the second argument. We experimented with two distinct styles for prompt design: Verbose Instructional Prompt and a Structured Templated Prompt. The verbose prompt, illustrated in Figure 1, uses natural language to explicitly define the model’s role, the task objective, the various input components, and decision-making guidelines. This contrasts with the structured prompt, which organizes the raw inputs into a compact, delimiter-separated format (e.g., ... \$\$ Arg1 \$\$ > ## Arg2 ## ...), resembling the input format for encoder models (Section 3.4). Although a large part of the verbose instructions is repeated identically in all samples, and may therefore be considered redundant, our experiments consistently showed that the verbose instructional prompts yielded superior performance, improving model accuracy by approximately 1–2% compared to the more concise, structured variants.

<sup>2</sup><https://huggingface.co/Qwen/Qwen3-4B>

<sup>3</sup><https://huggingface.co/JuNymphae/Georgetown-qwen3-4B-pruned-for-disrpt2025>



**Figure 1:** Illustration of the Verbose Instructional Prompt used in Qwen3-4B experiments.

### 3.4 mT5 Encoder

We also experimented with mT5 (Xue et al., 2021), a multilingual T5 variant pretrained on a Common Crawl-based corpus covering 101 languages. Specifically, we selected the mT5-XL variant, comprising 3.7 billion parameters, which comes closest to the shared task 4B limit. For our classification task, we used only the encoder and added a classification head. We explored two strategies for incorporating metadata and discourse features: (1) encoding them as special input tokens, and (2) using separate embedding layers concatenated with the encoder input.

**Feature Injection via Input Tokens** We prepended LCF features as special input tokens (e.g., LANG\_eng, FW\_erst, CORP\_gum), so that the model can incorporate this metadata directly in its tokenized input sequence. Since mT5 is trained to interpret prompt-like text, this approach naturally lets the model condition on task-specific context. This makes metadata injection especially effective in our setting, where the goal is to classify discourse relations across varied domains and annotation schemes.

In addition to metadata, we also applied this in-

jection strategy to categorical features from the DiscoDisco feature set (Gessler et al., 2021), which capture properties such as whether the units are full sentences, whether the relation is discontinuous, and whether the two units share the same speaker. We encoded these features as explicit key-value tokens, for example IS\_SENTENCE\_1, DISCONTINUOUS\_0, and SAME\_SPEAKER\_1. This design differs from the method of Metheniti et al. (2024), who append only raw feature values (e.g., 0.3, 0.5) in a fixed order, with each position implicitly corresponding to a particular feature. While their strategy reduces vocabulary size, it ties interpretation to positional indexing, making it less robust to reordering. By contrast, our approach makes the semantics of each feature explicit and order-independent, which aligns more naturally with mT5’s training paradigm and improves interpretability.

Furthermore, we implemented pseudo-directional features from DiscoDisco. Specifically, for relations labeled as left-to-right (1>2), we inserted direction using the tokens } and > before and after the first argument span to signal directional flow. For right-to-left (1<2) relations, the inverse

markers were used. These directional cues are lightweight but informative, and help disambiguate argument structure across instances, especially in genres with flexible syntax or conversational turn-taking.

The resulting input sequence is organized as follows: metadata (LCF features), followed by categorical DiscoDisco features, and finally the target argument span:

(3) LANG\_eng FW\_erst CORP\_gum [SEP]  
 IS\_SENTENCE\_1 DISCONTINUOUS\_0  
 SAME\_SPEAKER\_1 GENRE\_academic  
 [SEP] } Aesthetic Appreciation  
 and Spanish Art: > Arg2: In this  
 study we used eye-tracking in the  
 first stage

**Feature Embedding(s)** We hypothesized that treating argument spans (arg1 and arg2) together as a single sequence separated by a special token [SEP] might better leverage mT5’s native relative positional embeddings and attention dynamics. Meanwhile, surrounding context (Pre/Post) and metadata features were proposed to be embedded separately, since context can be long and sparse, which may attenuate positional signal strength if concatenated directly with argument spans. Thus, our proposed embedding schema is:

(4) Concat(Embed(meta + features) +  
 Embed(pre-context) + Embed(arg1 [SEP]  
 arg2) + Embed(post-context))

This setup preserves the association between argument spans, while mitigating positional confusion or noise from long context sequences.

Exploratory experiments on a smaller development subset indicated that this structure offered modest conceptual clarity but did not substantially outperform simply prepending LCF and categorical DiscoDisco features as special tokens.

## 4 Results

### 4.1 Performance Comparison of Encoder-Only and Decoder-Only Model

Table 3 reports the test scores of the encoder-only and decoder-only models across all 38 corpora.

We note several coarse observations at the outset: First, the decoder model outperforms the encoder model in all datasets except four: eng.dep.covdtb, eus.rst.ert, tur.pdtb.tdb, tha.pdtb.tdtb. However, the difference is minor (around 1 accuracy point or

Corpus	Decoder (DeDisCo)	Encoder
ces.rst.crdt	52.70	51.35
deu.pdtb.pcc	67.01	56.19
deu.rst.pcc	67.03	49.82
eng.dep.covdtb	68.21	73.05
eng.dep.scidtb	83.66	79.58
eng.erst.gentle	67.08	61.29
eng.erst.gum	73.45	62.98
eng.pdtb.gentle	67.94	61.07
eng.pdtb.gum	71.39	65.20
eng.pdtb.pdtb	83.77	77.32
eng.pdtb.tedm	71.79	61.54
eng.rst.oll	62.73	51.66
eng.rst.rstdt	73.27	62.60
eng.rst.sts	58.54	49.39
eng.rst.umuc	67.36	59.09
eng.sdrst.msdc	90.00	84.11
eng.sdrst.stac	75.80	65.96
eus.rst.ert	54.64	55.67
fas.rst.prstc	60.47	57.77
fra.sdrst.annodis	60.39	52.82
ita.pdtb.luna	70.13	66.13
nld.rst.nldt	68.62	53.85
pcm.pdtb.disconaija	59.39	41.40
pol.iso.pdc	74.02	55.05
por.pdtb.crpc	79.17	75.64
por.pdtb.tedm	68.41	64.84
por.rst.cstn	70.22	69.85
rus.rst.rst	74.85	68.95
spa.rst.rststb	69.72	64.55
spa.rst.sctb	83.02	76.73
tha.pdtb.tdtb	96.73	96.80
tur.pdtb.tdb	64.13	65.08
tur.pdtb.tedm	59.23	54.55
zho.dep.scidtb	80.00	68.37
zho.pdtb.cdtb	88.65	86.54
zho.pdtb.ted	75.49	66.24
zho.rst.gcdt	75.55	65.37
zho.rst.sctb	74.21	66.67
<b>Macro Average</b>	71.28	64.34
<b>Micro Average</b>	76.13	69.74

**Table 3:** Accuracy of encoder-only and decoder-only models on the test sets of 38 corpora. We use corpus codes for simplicity (check Table 6 in appendix D for language and framework information), with accuracy scores reported separately for the decoder-only and encoder-only models. Macro and micro averages are reported at the bottom.

less) in all but eng.dep.covdtb, which is a ‘test-only’ dataset, meaning the models have seen no data of the same kind – although a few other datasets are also test-only, they do have ‘related’ datasets (eng.erst.gentle, Aoyama et al. 2023, is modeled on eng.erst.gum, and the \*.tedm datasets closely follow recent versions of PDTB, Prasad et al. 2018). Aside from this unique property, we are unsure what sets this dataset apart.

By contrast, in the most extreme case the decoder achieves a 19 gain compared to the encoder on the Polish pol.iso.pdc. We can rule out data

contamination with gold data as a reason, since the data was released in the *surprise* test set and was reported to be annotated very recently.

Beyond individual dataset differences, we observed a broader trend with respect to data scale. In lower-resource settings, such as Czech (14.6K tokens), Dutch (24.9K), and French (32.7K), the decoder model consistently shows a substantial advantage over the encoder, with accuracy gains exceeding 10 points in some cases. These datasets have limited supervision and lack related training corpora, making them particularly reliant on pre-trained representations. The decoder’s autoregressive architecture and larger capacity appear to enable better generalization under these conditions. On the other hand, in larger datasets such as Thai (256K tokens), Turkish (496K), and English PDTB (1.17M), the performance gap narrows. In fact, the encoder slightly outperforms the decoder on Thai PDTB, suggesting that when sufficient labeled data is available, the simpler encoder-only setup can be just as effective, if not more so.

We also observed differences in how the two architectures respond to feature integration. LCF and DiscoDisco features consistently improved the encoder model, but in many cases degraded performance for the decoder. This suggests that encoder-only models can more effectively leverage categorical metadata and structural cues as additional signals, whereas decoder-only models are more sensitive to such injections. In contrast, extending the context window benefited the decoder but often harmed the encoder.

## 4.2 Decoder Model Ablation Tests

To assess the contribution of each additional features, we conduct an ablation study on the decoder-only model. In each experiment, one specific feature is removed from the input to evaluate its impact on performance. Detailed results are presented in appendix E (Table 7).

We find that direction is the most influential feature. When direction information is removed, 32 out of the 38 corpora experience a drop in accuracy of more than 4%, and 12 corpora suffer a decrease of over 10%. By contrast, for *tha.pdtb.tdtb*, direction contributes only marginal gains, and for *tur.pdtb.tedm*, it even leads to a performance drop. This is likely because Turkish is a free-word-order left-branching language, where in most subordinate constructions the clause with the connective precedes the main clause, causing direction fea-




Figure 2: Confusion matrix over the entire dataset.

tures based on linear order to introduce noise rather than useful signal (Zeyrek et al., 2009).

The second most impactful feature is context. Although prior work (Judge et al., 2024) suggests that adding context may not always be beneficial, we observe that context is a highly effective input for the decoder-only model. However, its impact varies significantly across corpora. For instance, *deu.rst.pcc* and *pcm.pdtb.disconaija* see gains exceeding 15%, and even up to 20% in *pol.iso.pdc*. For most corpora, the improvement is around 3%, though performance actually deteriorates on *tur.pdtb.tdb* and *ita.pdtb.luna*.

By comparison, LCF features tend to have smaller or even negative effects across many corpora. Nevertheless, they yield notable improvements (greater than 5%) for *deu.rst.pcc* and *zho.rst.sctb*.

A surprising finding emerges regarding data augmentation: it does not always improve performance on the target corpus. For example, For example, we observe gains on five corpora, except for *ces.rst.crdt* and *deu.pdtb.pcc*. On the other hand, for the source corpora, all five English datasets show consistent improvements from data augmentation, despite the source samples for the augmented entries also existing in the English training.

As for the DiscoDisco features, their overall impact was less pronounced across most corpora. The most notable exception was *eng.dep.covdtb*. For this unique, test-only dataset with no in-domain training data, removing these features surprisingly boosted performance by nearly 5%, suggesting they may introduce counterproductive noise.

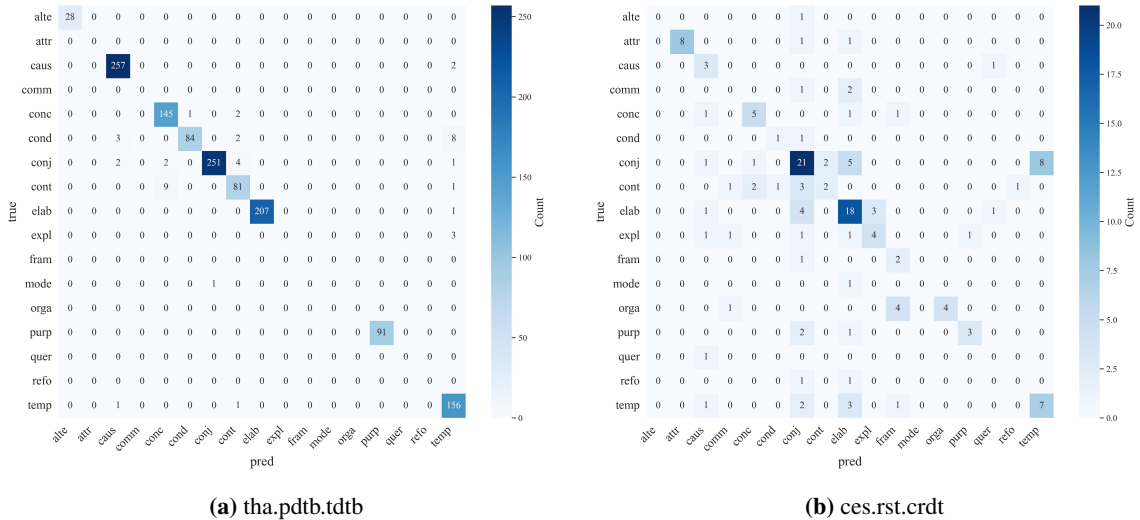


Figure 3: Confusion Matrices for Common Relations in the Highest and Lowest Scoring Datasets.

## 5 Error analysis

### Strengths and Confusions at the Relation Level

The confusion matrix in Figure 2 indicates very high performance on key discourse relations including ELABORATION, CONJUNCTION, CAUSAL, and TEMPORAL, demonstrating that the model reliably classifies these central categories. However, we observed minor yet notable confusion between ELABORATION and CONJUNCTION, and this might be because that ELABORATION and CONJUNCTION are frequently conflated due to their semantic and structural similarity, especially when explicit lexical signals are absent or ambiguous, which is often the case for ELABORATION, the most over-predicted label in the dataset. Teasing apart CONJUNCTION and TEMPORAL is also challenging, especially in cases where consecutive events are not signaled explicitly, but implicitly form a temporal succession relation.

### Strengths and Confusions at the Corpus Level

Figure 3 shows the confusion matrices for the highest- and lowest-scoring datasets, tha.pdtb.tdtb (TDTB) and ces.rst.crdt (CRDT). In CRDT, the model frequently defaults to the majority relation CONJUNCTION, reflecting strong over-prediction of common classes. The label distribution is also highly imbalanced: CONJUNCTION and ELABORATION dominate the dataset, while many relations (e.g., ALTERNATIVE, REFORMULATION, MODE) appear only rarely, making them difficult for the model to learn.

Relations marked by overt lexical cues, such as CONJUNCTION, achieve high accuracy in both datasets. Performance in TDTB is further aided by

the fact that this dataset covers only 12/17 possible labels, substantially reducing the possibilities for confusion compared to CRDT, which covers the full set of 17 labels.

## 6 Conclusion

We present DeDisCo, a decoder-only model, using a pruned Qwen3-4B basis, for the multilingual discourse relation classification task in the DISRPT 2025 Shared Task. Our system DeDisCo leverages supervised fine-tuning together with rich features, including metadata and instance-level cues such as unit distance, document position, and gold speaker information. We also incorporated augmented datasets to improve coverage for low-resource languages.

Our results suggest that decoder-only architectures are effective for this task, as their structure allows the model to integrate diverse sources of information (e.g., metadata, features, and context) within a unified text stream. Carefully designed instruction templates and feature injection further improve generalization, and natural prompt styles are helpful despite textual redundancy, even in a full fine-tuning setup. This enables the model to condition on a broad range of cues when making predictions, while adhering to a format the model is familiar with from its initial supervised fine-tuning. Error analysis highlights class imbalance as a persistent challenge, often leading to over-prediction of majority classes. Nonetheless, augmented data yielded measurable gains for low-resource languages, underscoring the importance of data enrichment strategies in this task.



## References

- Stergos Afantenos, Nicholas Asher, Farah Benamara, Myriam Bras, Cécile Fabre, Mai Ho-dac, Anne Le Draoulec, Philippe Muller, Marie-Paule Péry-Woodley, Laurent Prévot, Josette Rebeyrolles, Ludovic Tanguy, Marianne Vergez-Couret, and Laure Vieu. 2012. [An empirical resource for discovering cognitive principles of discourse organisation: the ANNODIS corpus](#). In *Proceedings of the Eighth International Conference on Language Resources and Evaluation (LREC'12)*, pages 2727–2734, Istanbul, Turkey. European Language Resources Association (ELRA).
- Kaveri Anuranjana. 2023. [DiscoFlan: Instruction fine-tuning and refined text generation for discourse relation label classification](#). In *Proceedings of the 3rd Shared Task on Discourse Relation Parsing and Treebanking (DISRPT 2023)*, pages 22–28, Toronto, Canada. The Association for Computational Linguistics.
- 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.
- Nicholas Asher, Julie Hunter, Mathieu Morey, Benamara Farah, and Stergos Afantenos. 2016. [Discourse structure and dialogue acts in multiparty dialogue: the STAC corpus](#). In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC'16)*, pages 2721–2727, Portorož, Slovenia. European Language Resources Association (ELRA).
- Peter Bourgonje and Manfred Stede. 2020. The pot-dam commentary corpus 2.2: Extending annotations for shallow discourse parsing. In *Proceedings of the 12th International Conference on Language Resources and Evaluation (LREC 2020) (to appear)*, Paris, France. European Language Resources Association (ELRA).
- 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.
- Chloé Braud, Amir Zeldes, Chuyuan Li, Yang Janet Liu, and Philippe Muller. 2025. [The DISRPT 2025 shared task on elementary discourse unit segmentation, connective detection, and relation classification](#). In *Proceedings of the 4th Shared Task on Discourse Relation Parsing and Treebanking (DISRPT 2025)*, pages 1–19, Suzhou, China. The Association for Computational Linguistics.
- Chloé Braud, Amir Zeldes, Laura Rivière, Yang Janet Liu, Philippe Muller, Damien Sileo, and Tatsuya Aoyama. 2024. [DISRPT: A multilingual, multi-domain, cross-framework benchmark for discourse processing](#). In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 4990–5005, Torino, Italia. ELRA and ICCL.
- Harry Bunt and Rashmi Prasad. 2016. ISO DR-Core (ISO 24617-8): Core concepts for the annotation of discourse relations. In *Proceedings 12th joint ACL-ISO workshop on interoperable semantic annotation (ISA-12)*, pages 45–54. Portoroz, Slovenia.
- Nicoletta Calzolari, Min-Yen Kan, Veronique Hoste, Alessandro Lenci, Sakriani Sakti, and Nianwen Xue, editors. 2024. *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*. ELRA and ICCL, Torino, Italy.
- Shuyuan Cao, Iria da Cunha, and Mikel Iruskietia. 2018. [The RST Spanish-Chinese treebank](#). In *Proceedings of the Joint Workshop on Linguistic Annotation, Multiword Expressions and Constructions (LAW-MWE-CxG-2018)*, pages 156–166, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- Paula Christina Figueira Cardoso, Erick Galani Maziero, Maria Lucía del Rosario Castro Jorge, M. Eloize, R. Kibar Aji Seno, Ariani Di Felippo, Lucia Helena Machado Rino, Maria das Graças Volpe Nunes, and Thiago Alexandre Salgueiro Pardo. 2011. CST-News - a discourse-annotated corpus for single and multi-document summarization of news texts in Brazilian Portuguese. In *Proceedings of the 3rd RST Brazilian Meeting*, pages 88–105, Cuiabá, Brazil.
- 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*.
- Yi Cheng and Sujian Li. 2019. [Zero-shot Chinese discourse dependency parsing via cross-lingual mapping](#). In *Proceedings of the 1st Workshop on Discourse Structure in Neural NLG*, pages 24–29, Tokyo, Japan. Association for Computational Linguistics.
- Alexis Conneau, Ruty Rinott, Guillaume Lample, Adina Williams, Samuel Bowman, Holger Schwenk, and Veselin Stoyanov. 2018. [XNLI: Evaluating cross-lingual sentence representations](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2475–2485, Brussels, Belgium. Association for Computational Linguistics.
- Iria da Cunha, Juan-Manuel Torres-Moreno, and Gerardo Sierra. 2011. [On the development of the RST Spanish treebank](#). In *Proceedings of the 5th Linguistic Annotation Workshop*, pages 1–10, Portland,

- Oregon, USA. Association for Computational Linguistics.
- Zeyu Dai and Ruihong Huang. 2018. [Improving implicit discourse relation classification by modeling inter-dependencies of discourse units in a paragraph](#). 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 141–151, New Orleans, Louisiana. Association for Computational Linguistics.
- Michel Génèreux, Iris Hendrickx, and Amália Mendes. 2012. [Introducing the reference corpus of contemporary Portuguese online](#). In *Proceedings of the Eighth International Conference on Language Resources and Evaluation (LREC'12)*, pages 2237–2244, Istanbul, Turkey. European Language Resources Association (ELRA).
- 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.
- Eduard H. Hovy. 1990. [Parsimonious and profligate approaches to the question of discourse structure relations](#). In *Proceedings of the Fifth International Workshop on Natural Language Generation*, Linden Hall Conference Center, Dawson, Pennsylvania. Association for Computational Linguistics.
- Julie Hunter, Hiroyoshi Yamasaki, Océane Granier, Jérôme Louradour, Roxane Bertrand, Kate Thompson, and Laurent Prévot. 2024. [MEETING: A corpus of French meeting-style conversations](#). In *Actes de la 31ème Conférence sur le Traitement Automatique des Langues Naturelles, volume 1 : articles longs et prises de position*, pages 508–529, Toulouse, France. ATALA and AFPC.
- Mikel Iruskieta, María Jesús Aranzabe, Arantza Diaz de Ilarraza, Itziar Gonzalez-Dios, Mikel Lersundi, and Oier Lopez de Lacalle. 2013. [The RST Basque TreeBank: An online search interface to check rhetorical relations](#). In *4th Workshop on RST and Discourse Studies*, pages 40–49, Fortaleza, Brasil.
- Yangfeng Ji and Jacob Eisenstein. 2014. [Representation learning for text-level discourse parsing](#). In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 13–24, Baltimore, Maryland. Association for Computational Linguistics.
- Evi Judge, Reece Suchocki, and Konner Syed. 2024. [An analysis of sentential neighbors in implicit discourse relation prediction](#). *Preprint*, arXiv:2405.09735.
- Alex Lascarides and Nicholas Asher. 2007. [Segmented discourse representation theory: Dynamic semantics with discourse structure](#). In *Computing meaning*, pages 87–124. Springer.
- Chuyuan Li, Chloé Braud, Maxime Amblard, and Giuseppe Carenini. 2024. [Discourse relation prediction and discourse parsing in dialogues with minimal supervision](#). In *Proceedings of the 5th Workshop on Computational Approaches to Discourse (CODI 2024)*, pages 161–176, St. Julians, Malta. Association for Computational Linguistics.
- Wei Liu, Yi Fan, and Michael Strube. 2023. [HITS at DISRPT 2023: Discourse segmentation, connective detection, and relation classification](#). In *Proceedings of the 3rd Shared Task on Discourse Relation Parsing and Treebanking (DISRPT 2023)*, pages 43–49, Toronto, Canada. The Association for Computational Linguistics.
- Wei Liu and Michael Strube. 2023. [Annotation-inspired implicit discourse relation classification with auxiliary discourse connective generation](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15696–15712, Toronto, Canada. Association for Computational Linguistics.
- Yang Janet Liu, Tatsuya Aoyama, Wesley Scivetti, Yilun Zhu, Shabnam Behzad, Lauren Elizabeth Levine, Jessica Lin, Devika Tiwari, and Amir Zeldes. 2024. [GDTB: Genre diverse data for English shallow discourse parsing across modalities, text types, and domains](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 12287–12303, 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.
- Daniel Marcu. 1999. [A decision-based approach to rhetorical parsing](#). In *Proceedings of the 37th Annual Meeting of the Association for Computational Linguistics*, pages 365–372, College Park, Maryland, USA. Association for Computational Linguistics.
- Daniel Marcu and William Wong. 2004. [A phrase-based, joint probability model for statistical machine translation](#). In *Proceedings of the Fourth International Conference on Language Resources and Evaluation (LREC)*, pages 219–222, Lisbon, Portugal. European Language Resources Association (ELRA).
- Xin Men, Mingyu Xu, Qingyu Zhang, Bingning Wang, Hongyu Lin, Yaojie Lu, Xianpei Han, and Weipeng Chen. 2024. [Shortgpt: Layers in large language models are more redundant than you expect](#). *Preprint*, arXiv:2403.03853.
- Amália Mendes and Pierre Lejeune. 2022. [Crpc-db a discourse bank for portuguese](#). In *Computational*

- Processing of the Portuguese Language: 15th International Conference, PROPOR 2022, Fortaleza, Brazil, March 21–23, 2022, Proceedings*, page 79–89, Berlin, Heidelberg. Springer-Verlag.
- Eleni Metheniti, Chloé Braud, and Philippe Muller. 2024. [Feature-augmented model for multilingual discourse relation classification](#). In *Proceedings of the 5th Workshop on Computational Approaches to Discourse (CODI 2024)*, pages 91–104, St. Julians, Malta. Association for Computational Linguistics.
- Eleni Metheniti, Chloé Braud, Philippe Muller, and Laura Rivière. 2023. [DisCut and DiscReT: MELODI at DISRPT 2023](#). In *Proceedings of the 3rd Shared Task on Discourse Relation Parsing and Treebanking (DISRPT 2023)*, pages 29–42, Toronto, Canada. The Association for Computational Linguistics.
- Noriki Nishida and Yuji Matsumoto. 2022. [Out-of-domain discourse dependency parsing via bootstrapping: An empirical analysis on its effectiveness and limitation](#). *Transactions of the Association for Computational Linguistics*, 10:127–144.
- Maciej Ogrodniczuk, Aleksandra Tomaszewska, Daniel Ziembicki, Sebastian Żurowski, Ryszard Tuora, and Aleksandra Zwierzchowska. 2024. [Polish Discourse Corpus \(PDC\): Corpus design, ISO-compliant annotation, data highlights, and parser development](#). In (Calzolari et al., 2024), pages 12829–12835.
- OpenAI, Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altmenschmidt, Sam Altman, Shyamal Anadkat, Red Avila, Igor Babuschkin, Suchir Balaji, Valerii Balcom, Paul Baltescu, Haiming Bao, Mohammad Bavarian, Jeff Belgum, and 262 others. 2024. [Gpt-4 technical report](#). *Preprint*, arXiv:2303.08774.
- Siyao Peng, Yang Janet Liu, and Amir Zeldes. 2022a. [Chinese Discourse Annotation Reference Manual](#). Research Report, Georgetown University (Washington, D.C.).
- Siyao Peng, Yang Janet Liu, and Amir Zeldes. 2022b. [GCDT: A Chinese RST Treebank for Multigenre and Multilingual Discourse Parsing](#). In *Proceedings of the 2nd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 12th International Joint Conference on Natural Language Processing*, pages 382–391, Online only. Association for Computational Linguistics.
- Lucie Poláková, Šárka Zikánová, Jiří Mírovský, and Eva Hajičová. 2023. [Czech RST discourse treebank 1.0](#). LINDAT/CLARIAH-CZ digital library at the Institute of Formal and Applied Linguistics (ÚFAL).
- Andrew Potter. 2008. Interactional coherence in asynchronous learning networks: A rhetorical approach. *Internet and Higher Education*, 11(2):87–97.
- Rashmi Prasad, Bonnie Webber, and Aravind Joshi. 2014. [Reflections on the Penn Discourse TreeBank, comparable corpora, and complementary annotation](#). *Computational Linguistics*, 40(4):921–950.
- Rashmi Prasad, Bonnie Webber, and Alan Lee. 2018. [Discourse annotation in the PDTB: The next generation](#). In *Proceedings of the 14th Joint ACL-ISO Workshop on Interoperable Semantic Annotation*, pages 87–97, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- Ponrawee Prasertsom, Apiwat Jaronpol, and Attapol T. Rutherford. 2024. [The thai discourse treebank: Annotating and classifying thai discourse connectives](#). *Transactions of the Association for Computational Linguistics*, 12:613–629.
- Laurent Prévot, Roxane Bertrand, and Julie Hunter. 2025. [Segmenting a large french meeting corpus into elementary discourse units](#). In *Proceedings of SIGDIAL*.
- Gisela Redeker, Ildikó Berzlánovich, Nynke van der Vliet, Gosse Bouma, and Markus Egg. 2012. [Multi-layer discourse annotation of a Dutch text corpus](#). In *Proceedings of the Eighth International Conference on Language Resources and Evaluation (LREC’12)*, pages 2820–2825, Istanbul, Turkey. European Language Resources Association (ELRA).
- Giuseppe Riccardi, Evgeny A. Stepanov, and Shammur Absar Chowdhury. 2016. [Discourse connective detection in spoken conversations](#). In *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6095–6099.
- Merel CJ Scholman, Marian Marchal, AriaRay Brown, and Vera Demberg. 2025. [Disconaija: a discourse-annotated parallel nigerian pidgin-english corpus](#). *Language Resources and Evaluation*, pages 1–37.
- Sara Shahmohammadi, Hadi Veisi, and Ali Darzi. 2021. [Persian Rhetorical Structure Theory](#). *arXiv preprint arXiv:2106.13833*.
- Manfred Stede, Stergos Afantenos, Andreas Peldszus, Nicholas Asher, and Jérémy Perret. 2016. [Parallel discourse annotations on a corpus of short texts](#). In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC’16)*, pages 1051–1058, Portorož, Slovenia. European Language Resources Association (ELRA).
- Manfred Stede and Arne Neumann. 2014. [Potsdam Commentary Corpus 2.0: Annotation for discourse research](#). In *Proceedings of the Language Resources and Evaluation Conference (LREC ’14)*, pages 925–929, Reykjavik.
- Kate Thompson, Julie Hunter, and Nicholas Asher. 2024. [Discourse structure for the Minecraft corpus](#). In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 4957–4967, Torino, Italia. ELRA and ICCL.

- Svetlana Toldova, Dina Pisarevskaya, Margarita Ananyeva, Maria Kobozeva, Alexander Nasedkin, Sofia Nikiforova, Irina Pavlova, and Alexey Shelepov. 2017. [Rhetorical relations markers in Russian RST treebank](#). In *Proceedings of the 6th Workshop on Recent Advances in RST and Related Formalisms*, pages 29–33, Santiago de Compostela, Spain. Association for Computational Linguistics.
- Aleksandra Tomaszewska, Purificação Silvano, António Leal, and Evelin Amorim. 2024. [ISO 24617-8 applied: Insights from multilingual discourse relations annotation in English, Polish, and Portuguese](#). In *Proceedings of the 20th Joint ACL - ISO Workshop on Interoperable Semantic Annotation @ LREC-COLING 2024*, pages 99–110, Torino, Italia. ELRA and ICCL.
- Sara Tonelli, Giuseppe Riccardi, Rashmi Prasad, and Aravind Joshi. 2010. [Annotation of discourse relations for conversational spoken dialogs](#). In *Proceedings of the Seventh International Conference on Language Resources and Evaluation (LREC'10)*, Valletta, Malta. European Language Resources Association (ELRA).
- Bonnie Webber, Rashmi Prasad, Alan Lee, and Aravind Joshi. 2019. The Penn Discourse TreeBank 3.0 annotation manual. Technical report, University of Edinburgh, Interactions, LLC, University of Pennsylvania.
- Linting Xue, Noah Constant, Adam Roberts, Mihir Kale, Rami Al-Rfou, Aditya Siddhant, Aditya Barua, and Colin Raffel. 2021. [mT5: A massively multilingual pre-trained text-to-text transformer](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 483–498, Online. Association for Computational Linguistics.
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, Chujie Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng Hu, Hao Ge, Haoran Wei, Huan Lin, Jialong Tang, and 41 others. 2025. [Qwen3 technical report](#). *Preprint*, arXiv:2505.09388.
- An Yang and Sujian Li. 2018. [SciDTB: Discourse dependency TreeBank for scientific abstracts](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 444–449, Melbourne, Australia. Association for Computational Linguistics.
- Cheng Yi, Li Sujian, and Li Yueyuan. 2021. [Unifying discourse resources with dependency framework](#). In *Proceedings of the 20th Chinese National Conference on Computational Linguistics*, pages 1058–1065, Huhhot, China. Chinese Information Processing Society of China.
- Yue Yu, Yilun Zhu, Yang Liu, Yan Liu, Siyao Peng, Mackenzie Gong, and Amir Zeldes. 2019. [GumDrop at the DISRPT2019 shared task: A model stacking approach to discourse unit segmentation and connective detection](#). In *Proceedings of the Workshop on Discourse Relation Parsing and Treebanking 2019*, pages 133–143, Minneapolis, MN. Association for Computational Linguistics.
- Karolina Zaczynska and Manfred Stede. 2024. [Rhetorical strategies in the UN security council: Rhetorical Structure Theory and conflicts](#). In *Proceedings of the 25th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 15–28, Kyoto, Japan. 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, Debopam Das, Erick Galani Maziero, Juliano Antonio, and Mikel Iruskieta. 2019. [The DISRPT 2019 shared task on elementary discourse unit segmentation and connective detection](#). In *Proceedings of the Workshop on Discourse Relation Parsing and Treebanking 2019*, pages 97–104, Minneapolis, MN. Association for Computational Linguistics.
- Amir Zeldes, Yang Janet Liu, Mikel Iruskieta, Philippe Muller, Chloé 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, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Deniz Zeyrek and Murathan Kurfalı. 2017. [TDB 1.1: Extensions on Turkish discourse bank](#). In *Proceedings of the 11th Linguistic Annotation Workshop*, pages 76–81, Valencia, Spain. Association for Computational Linguistics.
- Deniz Zeyrek, Amália Mendes, Yulia Grishina, Murathan Kurfalı, Samuel Gibbon, and Maciej Ogródniczuk. 2019. TED Multilingual Discourse Bank (TED-MDB): A parallel corpus annotated in the PDTB style. *Language Resources and Evaluation*, pages 1–27.
- Deniz Zeyrek, Amália Mendes, and Murathan Kurfalı. 2018. [Multilingual extension of PDTB-style annotation: The case of TED multilingual discourse bank](#). In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan. European Language Resources Association (ELRA).
- Deniz Zeyrek, Umit Deniz Turan, Cem Bozsahin, Ruket Cakici, Ayisigi B. Sevdik-Calli, Isin Demirsahin, Berfin Aktas, İhsan Yalcinkaya, and Hale Ogel. 2009. [Annotating subordinators in the Turkish discourse](#)

bank. In *Proceedings of the Third Linguistic Annotation Workshop (LAW III)*, pages 44–47, Suntec, Singapore. Association for Computational Linguistics.

Deniz Zeyrek and Bonnie Webber. 2008. A discourse resource for Turkish: Annotating discourse connectives in the METU corpus. In *Proceedings of the 6th Workshop on Asian Language Resources*.

Yuping Zhou, Jill Lu, Jennifer Zhang, and Nianwen Xue. 2014. Chinese Discourse Treebank 0.5 LDC2014T21.

## A Data

We train and evaluate our models using all the datasets provided by the shared task organizers.<sup>4</sup> In total, the benchmark is composed of 39 datasets, covering 13 languages and 6 frameworks. These datasets were obtained from the following corpora: the Czech RST Discourse Treebank 1.0 (Poláková et al., 2023), the Potsdam Commentary Corpus (Stede and Neumann, 2014; Bourgonje and Stede, 2020), the COVID-19 Discourse Dependency Treebank (Nishida and Matsumoto, 2022), the Discourse Dependency TreeBank for Scientific Abstracts (Yang and Li, 2018; Yi et al., 2021; Cheng and Li, 2019), the Genre Tests for Linguistic Evaluation corpus (Aoyama et al., 2023), the Georgetown University Multilayer corpus (Zeldes, 2017), the RST Discourse Treebank (Carlson et al., 2001), the Science, Technology, and Society corpus (Potter, 2008), the University of Potsdam Multilayer UNSC Corpus (Zaczynska and Stede, 2024), the Minecraft Structured Dialogue Corpus (Thompson et al., 2024), the Strategic Conversations corpus (Asher et al., 2016), the Basque RST Treebank (Iruskieta et al., 2013), the Persian RST Corpus (Shahmohammadi et al., 2021), the ANNOTation DIScursive corpus (Afantenos et al., 2012), the SUMM-RE corpus (Hunter et al., 2024; Prévot et al., 2025), the Dutch Discourse Treebank (Redeker et al., 2012), the Polish Discourse Corpus (Ogrodniczuk et al., 2024; Calzolari et al., 2024), the Cross-document Structure Theory News Corpus (Cardoso et al., 2011), the Russian RST Treebank (Toldova et al., 2017), the RST Spanish Treebank (da Cunha et al., 2011), the RST Spanish-Chinese Treebank (Cao et al., 2018), the Georgetown Chinese Discourse Treebank (Peng et al., 2022b,a), the DiscoNaija corpus (Scholman et al., 2025), the Penn Discourse Treebank (Prasad et al., 2014; Webber et al., 2019), the TED-Multilingual Discourse Bank (English)

<sup>4</sup><https://github.com/disrpt/sharedtask2025>

(Zeyrek et al., 2018, 2019), the LUNA Corpus Discourse Data Set (Tonelli et al., 2010; Riccardi et al., 2016), the Portuguese Discourse Bank (Mendes and Lejeune, 2022; Génèreux et al., 2012), the Thai Discourse Treebank (Prasertsom et al., 2024), the Turkish Discourse Bank (Zeyrek and Webber, 2008; Zeyrek and Kurfalı, 2017), and the Chinese Discourse Treebank (Zhou et al., 2014).

## B Feature Utilization in Decoder and Encoder Models

Feature Group	Decoder	Encoder
LCF features	+	+
DiscoDisco features	+	+
Direction	+	+
Context (window size)	+	–
Augmented dataset	+	+

**Table 4:** Feature groups and dataset augmentation used in decoder vs. encoder models (“+” = included, “–” = excluded).

DiscoDisco Feature	Decoder	Encoder
Genre	–	+
Children	–	–
Discontinuous	–	+
Is Sentence	–	+
Length Ratio	–	–
Same Speaker	+	+
Document Length	–	–
Position	+	–
Distance	+	–
Lexical Overlap	–	–

**Table 5:** Detailed inclusion (“+”) or exclusion (“–”) of all DiscoDisco features for each model.

## C Experimental Setup of Decoder

The model was trained on four NVIDIA H100 GPUs. Training one epoch took approximately three hours, and evaluation on the test sets required an additional one and a half hours. With a per-device batch size of 1, and using 16 gradient accumulation steps, the effective batch size was 64.

On rare occasions, the generative model produced outputs that were not part of the predefined set of valid labels. In such cases, our evaluation script replaced the output with a randomly selected valid label. This was extremely infrequent, occurring fewer than five times across all evaluations.

## D Corpora from DISRPT 2025 Shared Task

Language Name (Code)	Framework	Corpus LFC Short Code
Czech (ces)	RST	ces.rst.crdt
Standard German (deu)	PDTB	deu.pdtb.pcc
Standard German (deu)	RST	deu.rst.pcc
English (eng)	DEP	eng.dep.covdtb
English (eng)	DEP	eng.dep.scidtb
English (eng)	eRST	eng.erst.gentle
English (eng)	eRST	eng.erst.gum
English (eng)	PDTB	eng.pdtb.gentle
English (eng)	PDTB	eng.pdtb.gum
English (eng)	PDTB	eng.pdtb.pdtb
English (eng)	PDTB	eng.pdtb.tedm
English (eng)	RST	eng.rst.oll
English (eng)	RST	eng.rst.rstdt
English (eng)	RST	eng.rst.sts
English (eng)	RST	eng.rst.umuc
English (eng)	SDRT	eng.sdr.t.msdc
English (eng)	SDRT	eng.sdr.t.stac
Basque (eus)	RST	eus.rst.ert
Persian (fas)	RST	f as.rst.prstc
French (fra)	SDRT	fra.sdr.t.annodis
French (fra)	SDRT	fra.sdr.t.summre'
Italian (ita)	PDTB	ita.pdtb.luna
Dutch (nld)	RST	nld.rst.nldt
Nigerian Pidgin (pcm)	PDTB	pcm.pdtb.disconaija
Polish (pol)	ISO	pol.iso.pdc
Portuguese (por)	PDTB	por.pdtb.crpc
Portuguese (por)	PDTB	por.pdtb.tedm
Portuguese (por)	RST	por.rst.cstn
Russian (rus)	RST	rus.rst.rtt
Spanish (spa)	RST	spa.rst.rststb
Spanish (spa)	RST	spa.rst.sctb
Thai (tha)	PDTB	tha.pdtb.tdtb
Turkey (tur)	PDTB	tur.pdtb.tdb
Turkey (tur)	PDTB	tur.pdtb.tedm
Chinese (zho)	DEP	zho.dep.scidtb
Chinese (zho)	PDTB	zho.pdtb.cdtb
Chinese (zho)	PDTB	zho.pdtb.ted
Chinese (zho)	RST	zho.rst.gcdt
Chinese (zho)	RST	zho.rst.sctb

**Table 6:** Collation of the data in DISRPT 2025 Shared task with the corresponding Language and Framework for reference. The frameworks include RST (Mann and Thompson, 1988), PDTB (Marcu and Wong, 2004), DEP (Stede et al., 2016), SDRT (Lascarides and Asher, 2007), eRST (Zeldes et al., 2025), and ISO 24617-8 (Bunt and Prasad, 2016; Tomaszewska et al., 2024). The languages are sorted based on the language code (ISO 639-3<sup>5</sup>). The corpus fra.sdr.t.summre marked with ' is not part of the relation classification task.

<sup>5</sup><https://www.iso.org/iso-639-language-code>

## E Ablation Test Results of Decoder

Corpus	Decoder	w/o LCF		w/o DiscoDisco		w/o Direction		w/o Context		w/o Aug	
		<i>abs.</i>	<i>gain</i>	<i>abs.</i>	<i>gain</i>	<i>abs.</i>	<i>gain</i>	<i>abs.</i>	<i>gain</i>	<i>abs.</i>	<i>gain</i>
ces.rst.crdt	52.70	54.73	-2.03	55.41	-2.71	46.62	6.08	50.00	2.70	57.43	-4.73
deu.pdtb.pcc	67.01	63.92	3.09	65.98	1.03	62.37	4.64	62.37	4.64	69.07	-2.06
deu.rst.pcc	67.03	59.71	7.32	64.84	2.19	48.72	18.31	51.28	15.75	64.10	2.93
eng.dep.covdtb	68.21	70.46	-2.25	73.16	-4.95	57.39	10.82	71.04	-2.83	71.81	-3.6
eng.dep.scidtb	83.66	84.29	-0.63	83.87	-0.21	82.57	1.09	81.05	2.61	83.56	0.1
eng.erst.gentle	67.08	62.93	4.15	65.99	1.09	60.27	6.81	64.62	2.46	67.32	-0.24
eng.erst.gum	73.45	71.29	2.16	73.11	0.34	64.72	8.73	66.84	6.61	72.41	1.04
eng.pdtb.gentle	67.94	65.27	2.67	67.05	0.89	63.61	4.33	67.43	0.51	66.79	1.15
eng.pdtb.gum	71.39	68.00	3.39	70.86	0.53	65.38	6.01	68.47	2.92	70.8	0.59
eng.pdtb.pdtb	83.77	83.19	0.58	84.07	-0.30	73.22	10.55	82.61	1.16	83.80	-0.03
eng.pdtb.tedm	71.79	68.95	2.84	71.23	0.56	64.67	7.12	68.09	3.70	69.52	2.27
eng.rst.oll	62.73	60.89	1.84	60.52	2.21	49.45	13.28	59.41	3.32	59.04	3.69
eng.rst.rstdt	73.27	67.38	5.89	73.41	-0.14	68.54	4.73	69.51	3.76	72.99	0.28
eng.rst.sts	58.54	55.49	3.05	56.71	1.83	45.43	13.11	50.00	8.54	54.88	3.66
eng.rst.umuc	67.36	66.53	0.83	66.53	0.83	61.16	6.2	61.16	6.20	62.4	4.96
eng.sdr.msdc	90.00	89.75	0.25	90.03	-0.03	88.82	1.18	86.14	3.86	89.08	0.92
eng.sdr.stac	75.80	76.33	-0.53	75.98	-0.18	70.92	4.88	70.66	5.14	74.73	1.07
eus.rst.ert	54.64	56.49	-1.85	51.96	2.68	43.30	11.34	46.60	8.04	52.16	2.48
fas.rst.prstc	60.47	59.12	1.35	59.12	1.35	51.52	8.95	50.84	9.63	59.63	0.84
fra.sdr.annodis	60.39	56.04	4.35	61.19	-0.80	53.14	7.25	51.53	8.86	58.94	1.45
ita.pdtb.luna	70.13	70.40	-0.27	70.13	0	61.6	8.53	70.93	-0.80	72.53	-2.40
nld.rst.nldt	68.62	69.85	-1.23	69.54	-0.92	55.08	13.54	61.23	7.39	67.69	0.93
pcm.pdtb.disconaija	59.39	60.96	-1.57	61.16	-1.77	51.13	8.26	42.97	16.42	60.18	-0.79
pol.iso.pdc	74.02	73.08	0.94	73.62	0.40	62.99	11.03	53.97	20.05	72.14	1.88
por.pdtb.crpc	79.17	79.09	0.08	78.85	0.32	73.48	5.69	77.48	1.69	77.72	1.45
por.pdtb.tedm	68.41	68.68	-0.27	68.41	0	65.11	3.30	65.38	3.03	67.03	1.38
por.rst.cstn	70.22	70.96	-0.74	71.32	-1.10	69.12	1.10	70.59	-0.37	71.32	-1.1
rus.rst.rrt	74.85	74.81	0.04	75.31	-0.46	66.18	8.67	69.52	5.33	74.46	0.39
spa.rst.rststb	69.72	71.83	-2.11	70.66	-0.94	64.55	5.17	65.96	3.76	70.42	-0.70
spa.rst.sctb	83.02	77.99	5.03	80.50	2.52	70.44	12.58	76.73	6.29	86.16	-3.14
tha.pdtb.tdtb	96.73	96.88	-0.15	96.5	0.23	96.58	0.15	96.73	0	96.50	0.23
tur.pdtb.tdb	64.13	66.03	-1.90	66.75	-2.62	59.86	4.27	64.61	-0.48	66.75	-2.62
tur.pdtb.tedm	59.23	58.4	0.83	59.78	-0.55	59.5	-0.27	58.95	0.28	58.4	0.83
zho.dep.scidtb	80.00	78.6	1.40	77.21	2.79	69.77	10.23	74.42	5.58	76.28	3.72
zho.pdtb.cdtb	88.65	88.52	0.13	90.5	-1.85	83.91	4.74	87.34	1.31	88.79	-0.14
zho.pdtb.ted	75.49	76.09	-0.60	75.86	-0.37	67.97	7.52	71.95	3.54	75.79	-0.30
zho.rst.gcdt	75.55	73.66	1.89	75.13	0.42	62.96	12.59	70.51	5.04	76.71	-1.16
zho.rst.sctb	74.21	67.30	6.91	73.58	0.63	62.26	11.95	70.44	3.77	71.70	2.51
<b>Macro Average</b>	71.28	70.10	1.18	71.21	0.08	63.80	7.49	66.56	4.72	70.82	0.47
<b>Micro Average</b>	76.13	75.15	0.98	76.38	-0.25	69.53	6.60	72.24	3.89	75.86	0.27

**Table 7:** Accuracy results of the ablation study on the decoder-only model: next to the scores from Table 3, we report scores without LCF features, without DiscoDisco features, without direction, without context and without data augmentation, as well as the “gain” for each (non-ablated score – ablated score).