

An Exploratory Study of Discursive Signals Annotation in User-Generated Content

Rivânia da Paixão de Jesus Carvalho¹, Mateus Araújo Pereira¹, Jackson Wilke da Cruz Souza¹

¹Programa de Pós-Graduação em Língua e Cultura
Universidade Federal da Bahia - Salvador/Ba

rivaniacarvalho@bol.com.br {pereiramateus;jacksoncruz}@ufba.br

Abstract: *In this paper, we present the annotation of Discursive Signals (DSs) in User-Generated Content (UGC), based on a sample of tweets/posts from the financial-market domain annotated according to the Enhanced Rhetorical Structure Theory (eRST) model. The results reveal the presence of recurrent discourse relations in this type of text and their respective DSs, contributing to the expansion of the previously established taxonomy.*

Resumo: *Nesse artigo, apresentamos a anotação de Sinalizadores Discursivos (SDs) em Conteúdos Gerados por Usuários (UGC), com base em uma amostra de tweets/posts do domínio do mercado financeiro, anotados de acordo com o modelo da Enhanced Rhetorical Structure Theory (eRST). Os resultados revelam a presença de relações discursivas recorrentes nesse tipo de texto e seus respectivos SDs, contribuindo para a expansão da taxonomia previamente estabelecida.*

1. Introduction

Discourse analysis in Natural Language Processing (NLP) relies on theoretical models that derive discourse structure from semantic and syntactic information. One such model is Rhetorical Structure Theory (RST) [Mann and Thompson 1987], which views discourse as organized through coherence relations, allowing the examination of the communicative intentions of authors and possible interpretations derived by readers through the arrangement of Elementary Discourse Units (EDUs). In the context of the processing of Portuguese, discourse-level information plays a central role in tasks such as summarization, opinion mining, argumentation analysis, and dialogue systems [Taboada and Mann 2006].

To account for these phenomena, RST relations are often identified through Discursive Markers (DMs), typically conjunctions and prepositions. Although they can also be realized through syntactic, semantic, lexical, and morpholexical strategies, they are frequently inferred when no explicit markers are present. However, when dealing with less conventional text types, such as User-Generated Content (UGC) [Krumm, Davies, Narayanaswami 2008], RST faces challenges in the consistent identification of discourse relations. In response to these limitations, Zeldes *et al.* (2025) propose Enhanced Rhetorical Structure Theory (eRST), which extends RST by refining segmentation, introducing new relations, and expanding signaling strategies, offering a robust framework for representing discourse across different text types.

Figure 1¹ illustrates an example of discourse annotation based on the eRST framework. The text is segmented into six EDUs, connected through rhetorical relations that organize the informational structure of the discourse. EDUs 1 and 2 are linked by an *Elaboration* relation, in which the second EDU provides additional information about the announced plan. Similarly, EDU 3 is connected to EDUs 4 and 5 through a *List* relation, signaling the enumeration of investment targets. At a higher level, the discourse structure connects the previous segments to EDU 6 through an *Evaluation* relation, expressing the author's interpretation of the reported information.




Figure 1 - Discourse Annotation.

The annotated example also presents different DSs that support the interpretation of rhetorical relations in the text. In the *Elaboration* relation between EDUs 2-4 and 3-4, the nominal complement structure “*plan of an incentive*” functions as a *lexical-syntactic* cue. In the *joint-list* relation connecting EDUs 3 and 4, the conjunction “*and*” signals enumeration, while structural parallelism reinforces the interpretation of multiple investment targets. *Evaluation* relation associated with EDU 5 is supported by *evaluative lexical* choices and by the hashtag, which functions as a genre-specific cue frequently observed in financial-domain UGC texts. These elements illustrate how rhetorical relations in UGC may be licensed by lexical, syntactic, structural, and platform-specific cues.

Das and Taboada (2018) argue that DSs must be present whenever a rhetorical relation can be inferred, noting that most relations rely not on explicit DMs, but on structural and linguistic cues. For English, the authors propose a DS taxonomy based on corpora previously annotated by themselves. A DS taxonomy for Brazilian Portuguese was proposed by Dantas *et al.* (2024) developed a DS taxonomy based on the analysis of the CSTNews corpus [Cardoso *et al.* 2011], covering Graphical, Semantic, Syntactic, Morphological, and DM cues. This taxonomy was later expanded by Gama and Souza (2024) to better account for UGC, with the introduction of the Genre category, which deals with elements such as emojis and hashtags.

Figure 2 presents the DS taxonomy proposed for Brazilian Portuguese and its extension to UGC texts. The taxonomy organizes signals into six major categories: *Graphical*, *Semantic*, *Syntactic*, *Morphological*, *Discourse Marker*, and *Genre*. While the first categories capture traditional linguistic signaling strategies, the *Genre* category was introduced to account for platform-specific elements frequently observed in UGC, such

¹ Original tweet: China anunciou um plano de incentivo à economia. O governo chinês pretende aumentar investimentos em ferrovias e moradias. Bom p/ #VALE5

as *hyperlinks*, *retweets*, *hashtags*, and *emojis*. This taxonomical organization supports the identification of rhetorical relations beyond prototypical DMs.




Figure 2 - Discourse Signals translated from Gama and Souza (2024).

Using UGC texts from the financial domain annotated according to the eRST model, researchers conducted both manual and automatic DS identification. In this process, LLMs were employed to detect potential DSs, and their outputs were compared with human annotations, following procedures described in Authors Pereira and Souza (2025). This process resulted in an enhanced *corpus* with an additional annotation layer and contributed to the refinement of the DS taxonomy.

The aim of this study is to identify DSs that can serve as reliable cues for the characterization of eRST relations between EDUs in the financial market tweets/posts domain. To achieve this, we first conducted a manual annotation of EDUs, eRST relations, and their associated DSs in a sample of the *corpus*. Based on this annotation, we describe patterns of association between types of DSs and rhetorical relations in UGC.

A comparative analysis was made between the manual annotation approach and the use of LLMs for the automatic identification of EDUs, eRST relations, and DSs. Subsequently, we examine points of convergence and divergence between the two approaches, as well as their respective limitations. To this end, we compare the model’s performance with human annotations and discuss the potential implications for computational discourse modeling.

This ongoing study contributes to this effort by exploring how DSs can be systematically identified, categorized, and partially automated within the eRST framework, with a particular focus on Portuguese tweets/posts from the financial domain. This paper is organized into five sections beyond this Introduction. In Section 2, we present some related works. In Section 3, we outline the methodology employed. Section 4 discusses the results and relevant insights, and in Section 5, we offer concluding remarks.

2. Related work

In recent years, there has been a growing interest in refining LLMs and adapting them to new text genres, particularly UGC, whose informal, fragmented, and platform-dependent nature increases analytical complexity. UGC texts, particularly those published on social

media platforms, exhibit distinctive linguistic features that support the idea that coherence does not depend solely on prototypical DMs but rather on discourse structure [Zeldes *et al.* 2025]. Thus, both the textual properties of the genre and the discursive choices of the interlocutors are shaped by the constraints and affordances of each social media platform.

Launched in 2006, Twitter (now known as “X”) has become a significant public forum and has drawn increasing interest from scholars with the aim of understanding user behavior through linguistic manifestations. The texts produced on Twitter platform consist of content shared by individuals or institutions, based on real experiences, evaluations, ideas, comments, and opinions about events, brands, services, or products. According to Santos (2022), these texts can be considered UGC as they are produced and published by users on independent channels of online digital platforms, serving communicative purposes, either individually or in conjunction with other contributions from the same source or other sources.

Di Felippo *et al.* (2021) highlight the unique characteristics of UGC texts on Twitter/X, shaped by platform-specific conventions and limitations, and their implications for automated language processing. Given the informal structure of tweets/posts, the authors identified structural, orthographic and lexical features of UGC texts to define the appropriate unit of analysis.

These characteristics were systematized by Di Felippo *et al.* (2021) into seven dimensions, which include the following phenomena: *code simplification*, such as the addition or substitution of graphic signs like diacritics and hyphens; *abbreviation*, frequently occurring in the form of acronyms; *expression of feelings*, which, most of the time, appear in the form of emoticons and repetition of punctuation; *influence of foreign languages*, especially English; *expression of orality*, which reflects colloquial communication; *metalinguistic elements*, which generally appear at the end of the tweet and are tokenized together with the word or expression they precede; and the *phenomenon of dominance*, which distinguishes DANTEStocks [Di Felippo *et al.*, 2021] tweets from other tweets/posts.

As eRST aims to offer a more comprehensive representation of rhetorical relations in texts of any genre, as asserted by Zeldes *et al.* (2025), it significantly contributes to the computational modeling of discourse, integrating different levels of linguistic and contextual analysis. This integration has helped to identify real usage patterns in UGC texts. In addition to organizing the relations into a taxonomy adapted to the textual genre, this more comprehensive proposal allows for the characterization of explicit and implicit linguistic signals realized in the text. In line with the studies of [Zeldes *et al.* 2025] and [Di Felippo *et al.* 2021], this research seeks to provide the conditions necessary to capture the complexity, subjectivity, and hybridity typical of UGC texts, based on a taxonomy that enables the investigation and validation of new discursive relations guided by empirical evidence.

In Souza, Cardoso, and Rodrigues (2025) used GPT-4 and LLaMA-4 to identify DMs in news texts, revealing points of convergence and divergence both between the models and in comparison, with human annotation. In the RST relations of Attribution, Elaboration, and Circumstance, LLaMA tends to produce more concise annotations, often reduced to a single central DS such as communication verb *Attribution* or relative clause *Elaboration*. In such cases, GPT-4 offers more detailed and granular descriptions, sometimes excessively, by including additional Graphical, Semantic, or Morphological features. In several instances, both LLMs proposed DSs not included in the existing

taxonomy, such as Subordinate clause, Subordinating conjunction, and Referential element. This indicates potential directions for expanding or revising the typology and highlighting challenges in consistently applying the current labels.

The analysis also revealed certain limitations, including hallucinations and an excessive tendency toward simplification in both GPT-4 and LLaMA, as well as difficulties in identifying DSs in specific EDUs. The results show that the combination of manual and automatic annotation is productive in revising, consolidating, and potentially improving the taxonomy. The study further underscores the need to examine less frequent RST relations, expand the scope to inter-sentential relations, and explore prompt adjustments and hyperparameter tuning. In general, the comparison between models and human annotation highlights both the potential of LLMs for DS annotation and the particularities and limits of their information generation. Among these peculiarities, regarding UGC texts generated on Twitter, in addition to rhetorical relations, we observe DSs with characteristics specific to the genre, as noted by [Zeldes *et al.* 2025].

3. Materials and methods

To perform a comparative analysis between human annotation and the LLMs' performance, in this study, a pilot subset of the DANTEStocks *corpus* [Di Felippo *et al.* 2021] was used, consisting of 60 tweets written in Brazilian Portuguese and collected from the X/twitter platform. Due to the structural variability of the texts, to better account for UGC, the dataset was classified into three categories: Well-structured (WS), Moderately Structured (MS), and Poorly Structured (PS), following the framework proposed by Authors Pereira and Souza (2024).

This classification supports a more precise identification of discourse relations by addressing the distinct phenomena associated with each text type. Table 1 summarizes the criteria adopted for the classification of texts into the WS, MS, and PS categories. This categorization contributes to discourse annotation tasks, since structural variation in UGC directly affects EDU segmentation and rhetorical relation identification.

The subset had previously been manually reviewed and annotated by the Authors Pereira and Souza (2025) using the eRST framework proposed by Amir Zeldes (2025). The manual annotation was carried out by four annotators with academic training in NLP and discourse analysis. Annotation sessions were conducted synchronously and collaboratively, using the rstWeb tool [Zeldes 2016]. Each tweet was annotated by at least two annotators, with disagreements discussed during follow-up sessions.

This is an initial exploratory study whose pilot sample does not yet have quantitative metrics; therefore, the values found are presented in a qualitative and interpretive way, derived directly from the authors' analytical description. However, formal inter-annotator agreement metrics are still under computation. At this stage, the annotation process should be regarded as exploratory, aimed at refining guidelines and identifying ambiguous cases rather than producing a finalized gold standard. The annotator group was guided by the taxonomy and definitions of each type and subtype of DSs proposed by the authors Gama and Souza (2024) and Pereira and Souza (2024).

In addition, as a complementary approach and based on Souza, Cardoso and Rodrigues (2025), we experimented with automatic annotation using a GPT-4 as the LLM. The model was prompted in a few-shot setting to identify: (i) eRST relations, (ii) EDUs, and (iii) associated DSs, based on the taxonomy proposed in earlier work. We

emphasize that the LLM-based annotation is treated as a baseline, not as a replacement for human annotation.

4. Results and Discussion

Based on the DS taxonomy proposed by the Authors Gama and Souza (2024), we identified a range of linguistic, structural, and genre-specific features associated with eRST relations. According to the previous findings of the Authors Pereira and Souza (2025), the most frequent relations observed in this pilot sample were *Attribution-positive*, *Elaboration-additional*, *Joint-list*, and *Organizational-preparation*. This result indicates the prominence of source attribution, informational expansion, and organizational strategies in these types of UGC texts.

In Figure 3, we present examples that compare manual and automatic annotations. These examples are organized according to Pereira and Souza (2024), representing WS, MS, and PS, respectively. Based on these examples, we observe how both approaches (manual and automatic) behave in generating annotation results. At first glance at the examples, variation can be observed in the number of EDUs identified and in discourse segmentation; the manual annotation identified four EDUs, whereas the automatic annotation identified only two. Therefore, they present divergences in the attribution of rhetorical relations and DSs.

In some cases, the same DSs were interpreted as licensing different relations depending on the annotation method, underscoring the inherent ambiguity of DSs in UGC and the sensitivity of annotation outcomes to methodological choices.

A more detailed analysis reveals that these differences are conditioned by the level of *textual structure*. In the WS tweet, manual and automatic annotations show alignment, although the automatic approach simplifies the discourse structure to identify fewer relations and signals. For example, in (3) some eRST relations are preserved across both approaches (e.g., *Organization-preparation* and *Elaboration-additional*); and some relations are identified in a way that partially resembles the manual analysis.

In the MS tweet, we observed partial convergence between both approaches. Differences emerge in the interpretation of relations and in the attribution of DSs, although there is some overlap in the identification of specific signals, such as “RT”, which are identified in both annotations. In the PS tweet, divergences become more pronounced, with substantial variation across both prompted annotation tasks, as well as in the attribution of DSs.

Manual analysis confirms that the eRST framework can be consistently applied to tweets from the financial domain, despite their brevity and informal nature. The recurrence of a small set of relations suggests that beneath the apparent fragmentation of tweets, relatively stable discourse patterns contribute to textual coherence. These patterns are particularly evident in WS tweets, where organizational and elaborative relations are systematically licensed by explicit cues.

Tweets	MANUAL ANNOTATION				AUTOMATIC ANNOTATION			
	RST relationship	Signals			RST relationship	Signals		
		Mark	Type	Subtype		Mark	Type	Subtype
(1) 04/25 The Ilovespa closed down 0.81%, at 51,399 points. Highlights: AEDU3 - 3.72%, MFRG3 - 2.86%, and HGTX - 2.74%.	I. context-circumstance	"04/25"	DM	Adverb	I. organization-preparation	"Highlights"	Lexical	Textual organization
	II. elaboration-additional	"closed down"	Semantics	Verbal sense	II. elaboration-additional	---	---	---
	III. organization-preparation	:	Grafic	Punctuation	---	---	---	---
	IV. joint-list	,	Grafic	Punctuation	---	---	---	---
(2) RT @user: Joana is kind of addicted to a game	attribution-positive	RT / @	Text genre	Source	organization-phatic	RT	Pragmatics	Repost
(3) RT @usuário: Look at #PETRA there - Read 'Petrobras debt'...	I. attribution-positive	RT / @	Text genre	Source	organization-phatic	"Look"	---	---
	II. elaboration-additional	"Read"	Semantics	Verbal sense				

Table 1 - Manual and automatic analyses².

Regarding discursive signaling, the results demonstrate that eRST relations in UGC are rarely exclusively marked by prototypical DMs. *Attribution-positive* relations, for example, are frequently signaled by platform-specific elements such as retweet markers, user mentions, and explicit references to information sources, although these cues do not necessarily co-occur, especially in poorly structured tweets. *Elaboration-additional* relations are predominantly associated with hyperlinks and named entities (most notably hashtags), which function both as mechanisms of informational expansion and as tools for thematic indexing within the platform.

An important difference between manual and automatic annotation concerns the treatment of DSs and their classification. While manual annotation consistently applies the Type and Subtype distinctions defined in the DS taxonomy, the automatic annotation does not explicitly account for these categories. Nevertheless, the signals identified by LLM provide valuable empirical evidence that can inform future refinements of DSs taxonomy from a computational perspective. In Figure 1, this difference is particularly evident in the automatic annotation to capture signals at a more general level, whereas manual annotation provides a more fine-grained categorization.

Despite the divergences observed between manual and automatic annotation, the results suggest that LLMs approaches can contribute to identifying new possibilities for the classification of signals. As illustrated in Figure 1, the automatic annotation introduces signal types as *lexical* cues, which are not included in DS taxonomies. This evidence suggests that LLMs can reveal patterns that are not captured by current theoretical frameworks.

Analysis of the automatically annotated data further reveals a gradient of discourse organization across the three tweet categories. WS tweets rely primarily on *lexical* and *organizational signals*, such as heading-like expressions (e.g., "Highlights:"), *colon introducers*, *conditional complementizers*, and *quantificational expressions*. These signals systematically license relations such as *organization-preparation*, *joint-list*,

² Original tweets: (1) 25/04 O Ibovespa fechou com queda de 0,81%, aos 51.399 pontos. Destaques: AEDU3 (+3,72%), MFRG3 (+2,86%) e HGTX (+2,74%). | (2) RT @usuário: Joana é meio viciada em um jogo. | (3) RT @usuário: Olha a #PETR4 aí. Leia dívida da Petrobras...

elaboration-additional, and *contingency-condition*, yielding relatively explicit nucleus–satellite configurations and a reduced reliance on contextual inference.

MS tweets exhibit a hybrid signaling strategy that combines lexical cues with pragmatic, typographic, and interactional signals (e.g., *quotation marks* and *retweets*). In contrast, PS tweets show a marked decrease in canonical lexical signals and a sharp increase in platform-specific and paralinguistic cues (e.g., *user mentions*, *hashtags*, *emojis*, and *expressive punctuation*). In these cases, discourse coherence is predominantly sustained through organization-phatic and evaluation-comment relations, highlighting the central role of interactional conventions in UGC.

Overall, the comparison between manual and automatic annotation reveals three recurring tendencies: (i) LLMs tend to over-segment EDUs in densely packed or syntactically compressed tweets; (ii) rhetorical relations are often identified correctly at a coarse level but misclassified at finer-grained distinctions; and (iii) discursive signals are frequently detected yet inconsistently mapped onto the predefined Type and Subtype taxonomy. Taking together, these observations suggest that current LLMs implicitly capture discourse-relevant information but lack explicit control over linguistically motivated categories.




Figure 3 - Quantitative distribution of RST relations on DANTE corpus.

5. Final remarks

In this work, the annotation of DSs in Portuguese UGC texts was investigated based on a pilot sample of financial-domain tweets. The application of the eRST framework made it possible to identify recurrent discourse relations and their associated signals, even in a textual genre characterized by brevity, informality, and strong platform constraints. The results indicate that discourse coherence in UGC is sustained by a heterogeneous set of linguistic, structural, and platform-specific cues, whose distribution varies according to the degree of discourse organization of the texts.

From a computational perspective, the exploratory comparison between manual annotation and LLM-based automatic annotation reveals both opportunities and limitations. Although the LLM is often able to identify candidates DSs and rhetorical relations, it lacks systematic control over segmentation decisions and over the mapping of signals to linguistically motivated taxonomies. These findings suggest that LLMs are better viewed, at this stage, as semi-automatic support tools rather than as standalone solutions for discourse annotation.

Furthermore, despite the divergences observed in automatic annotation, the results indicate its potential to suggest new ways of classifying DSs. These results suggest categories that are not represented in existing taxonomies. Rather than being viewed only as a source of inconsistency, it can be considered a valuable source for generating new taxonomic categories for DS classification. This can contribute to future research supporting the expansion and refinement of the classification framework.

The relatively small sample size restricts the generalization of the findings and precludes a robust quantitative evaluation. Therefore, no hyperparameter adjustments were made as this is an initial exploratory study. In addition, the current DSs taxonomy does not yet account for more complex structural properties of discourse, such as hierarchical depth and inter-relational dependencies within the discourse tree.

In future work, we will focus on expanding the *corpus* to include other UGC texts, improving the prompt, refining the DSs taxonomy, and formalizing annotation guidelines. We also plan to scale the annotation process through hybrid human-AI workflows and to introduce standard evaluation metrics for automatic methods. Another promising direction is to compare the distribution of rhetorical relations identified in financial tweets with those observed in other textual genres and domains, to investigate whether this type of UGC presents distinctive discourse organization patterns.

In addition, future research will examine whether providing large language models with texts that have already been segmented into EDUs results in rhetorical analyses that more closely align with human annotations, potentially improving the quality of automatic discourse parsing. This line of research aims to contribute linguistically grounded resources and methodologies that support discourse-aware NLP applications for Portuguese in real-world digital environments.

Acknowledgments

This work was carried out at the Center for Artificial Intelligence of the University of São Paulo (C4AI – <http://c4ai.inova.usp.br/>), with the support of the São Paulo Research Foundation (FAPESP, Grant No. 2019/07665-4) and IBM Corporation. The project was also supported by the Brazilian Ministry of Science, Technology and Innovation (MCTI), through funding provided under Law No. 8,248 of October 23, 1991, within the scope of the PPI-SOFTEX Program, coordinated by Softex and published as ICT Residency 13 (Official Gazette reference: DOU 01245.010222/2022-44). In addition, this work received financial support from the FAPESB/CNPq Call No. 004/2023 – First Projects Program.

References

- Cardoso, P. C., Maziero, E. G., Jorge, M. L. C., Seno, E. M., Di Felippo, A., Rino, L. H. M., Nunes, M. d. G. V., and Pardo, T. A. (2011). Cstnews-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. sn.
- Dantas, E., Bárbara, L. de J. S., Pereira, M. A., Gama, N. S., Almeida, T. J. A., Souza, J. W. da C., et al. (2024). Manual de anotação de sinalizadores discursivos em textos jornalísticos. São Carlos: Instituto de Ciências Matemáticas e de Computação, Universidade de São Paulo. Recuperado de <https://repositorio.usp.br/directbitstream/31779c72-b91a-4f9d-8f40-799357711178/3207370.pdf>

- Das, D. and Taboada, M. (2018). Signalling of coherence relations in discourse, beyond discourse markers. *Discourse Processes*, 55(8):743–770.
- Di Felippo, A., Postali, C., Ceregatto, G., Gazana, L. S., da Silva, E. H., Roman, N. T., and Pardo, T. A. (2021). Descrição preliminar do corpus dantestocks: Diretrizes de segmentação para anotação segundo Universal Dependencies. In *Simpósio Brasileiro de Tecnologia da Informação e da Linguagem Humana (STIL)*, pages 335–343. SBC.
- Gama, N., & Souza, J. (2024). *Estudo preliminar sobre sinalizadores discursivos para Conteúdo Gerado por Usuário*. In *Anais do XV Simpósio Brasileiro de Tecnologia da Informação e da Linguagem Humana*, (pp. 418-423). Porto Alegre: SBC. doi:10.5753/stil.2024.245060
- Mann, W. C. and Thompson, S. A. (1987). Rhetorical structure theory: A framework for the analysis of texts. Technical report.
- Narayanaswami, C. (2008). User-generated content. *IEEE Pervasive Computing*.
- Pereira, M., & Souza, J. (2024). Subsídios *Linguísticos para Classificação Automática de Textos de User-Generated Content*. In *Anais do XV Simpósio Brasileiro de Tecnologia da Informação e da Linguagem Humana*, (pp. 429-433). Porto Alegre: SBC. doi:10.5753/stil.2024.245132
- Pereira, M., & Souza, J. (2025). Anotação Enhanced Rhetorical Structure Theory em textos de User-Generated Content. In *Anais do XVI Simpósio Brasileiro de Tecnologia da Informação e da Linguagem Humana*, (pp. 713-717). Porto Alegre: SBC. doi:10.5753/stil.2025.37877
- Santos, M. L. B. d. (2022). The “so-called” ugc: an updated definition of user-generated content in the age of social media. *Online Information Review*, 46(1):95–113.
- Souza, J. W. da C., Cardoso, P. C. F., & Rodrigues, R. (2025). Taxonomy of discourse signals for RST relations: A study in a news corpus. *Corpus Linguistics: Studies and Applications*, 25(1), e49793. <https://doi.org/10.1590/1984-6398202549793>
- Taboada, M. and Mann, W. C. (2006). Applications of rhetorical structure theory. *Discourse studies*, 8(4):567–588.
- Zeldes, A. (2016). rstweb-a browser-based annotation interface for rhetorical structure theory and discourse relations. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Demonstrations*, pages 1–5.
- Zeldes, A., Aoyama, T., Liu, Y. J., Peng, S., Das, D., and Gessler, L. (2025). erst: A signaled graph theory of discourse relations and organization. *Computational Linguistics*, 51(1):23–72.