

DM Doppelgangers: Implicit Connectives as eRST Signals

Lin Ai and Amir Zeldes
Corpling Lab
Georgetown University
{la987, amir.zeldes}@georgetown.edu

Abstract

Recent work representing discourse relations such as *cause* or *concession* in the framework of eRST has connected hierarchical discourse parsing to explicit connectives, such as ‘because’ or ‘although’, bringing the framework closer to lexicalized shallow parsing in the tradition of PDTB. However, while PDTB postulates implicit, unexpressed connectives (i.e. an implied *although* etc.), no such devices are recognized in eRST, and consequently next to nothing is known about the relationship between PDTB-style implicit connectives and eRST-style discourse graphs. In this paper we propose and evaluate an algorithm to align eRST data, which already indicates explicit connectives, to implicit connective annotations following the PDTB guidelines. We also conduct the first evaluation of the relationship between hierarchical RST-style relations and PDTB implicit connectives.

1 Introduction

Discourse is not a sequence of random, unrelated utterances that can be understood in isolation (Stede, 2012). Instead, one proposition can elaborate on another, contrast with it, or provide an explanation for it, as indicated in (1) by the choice between the explicit connectives *and/although/because* respectively. While a range of frameworks have been proposed to describe such relations, an overarching distinction can be made between, on the one hand, ‘shallow’ frameworks such as the Penn Discourse Treebank (PDTB, Prasad et al. 2014), which mainly seek to identify relations and the lexical means of signaling them (e.g. connectives), and on the other, ‘deep’ frameworks, such as Rhetorical Structure Theory (RST, Mann and Thompson 1988), which assume that relations form hierarchical structures, as in (2).

- (1) [Kim came]₁ [**and/although/because** her sister didn’t.]₂

- (2) [Kim came.]₁ [Her sister didn’t come.]₂ [I asked her to do that.]₃

The propositions in (2) are ambiguous in two ways: 1. they lack explicit connectives, meaning we can imagine, for example, an implicit *although* between units 1 and 2; and 2. they can be construed hierarchically as (1 (2 3)) if we assume for example that the sister did not come because I asked; or as ((1 2) 3) if we assume that Kim came, although her sister did not, and Kim did so because I asked. An RST tree fragment illustrating the second analysis is shown in Figure 1.




Figure 1: RST tree for the second reading of (2), visualized using rstWeb (Gessler et al., 2019).

The relationship between these two types of implicit information – implicit connectives and hierarchy – is understudied, due to the fact that the main source of data for studying implicit connectives comes from datasets following the PDTB framework, while the main source for studying discourse relation hierarchies is RST data, which does not recognize implicit connectives in its underlying theory. The main goal of this paper is to close this gap, by exhaustively associating relations in an RST-style paradigm with implicit connectives as defined in PDTB, thereby allowing us to study the interrelations between the two, and, as a byproduct of considering implicit connectives in discourse graphs, to identify additional relations which we would otherwise miss.

The main contributions of this paper are:

- introducing the notion of implicit connectives to RST style annotation
- release of the first dataset of enhanced RST relations aligned to implicit connectives
- expanding the relations annotated in the data with additional edges based on implicit connectives, thereby subsuming more of what PDTB-style annotation covers
- present first results on the interactions between relations, implicit connectives and other types of signals

2 Related work

2.1 PDTB and implicit connectives

The Penn Discourse TreeBank (PDTB) is a shallow discourse annotation framework developed using the English Wall Street Journal corpus (Marcus et al., 1993). Unlike tree-based discourse formalisms, PDTB does not assume that discourse relations form hierarchical structures. Instead, it annotates relations between pairs of text spans, typically signaled by lexical items such as connectives, including subordinating or coordinating conjunctions, discourse adverbials (e.g. ‘then’) and prepositional phrases (e.g. ‘in fact’). In addition, PDTB allows for implicit relations. These are the focus of the present paper and are identified based on whether an implicit connective could plausibly be inserted between two adjacent segments, such as an implied ‘although’ between units 1 and 2 in (2) above. Implicit connectives may co-occur with explicit discourse adverbials when they encode discourse information that is not fully conveyed by the overt adverbials. In an example by Prasad et al. (2014), although (3) and (4) are both signaled by ‘instead’, they express different discourse relations. In (3), we can naturally insert “but”, whereas in (4), “so” fits better. Implicit connectives are therefore “doppelgangers” of explicit dms in that they, too, can encode discourse relations.

- (3) Logically, she should be dead. (but) Instead, she feels fine. (substitution and concession)
- (4) He suspected he shouldn’t say that. (so) Instead he lied. (substitution and result)

2.2 RST, signals and eRST

RST (Mann and Thompson, 1988) is a pragmatic theory that assigns discourse relations based on the writer’s or speaker’s communicative intent. The minimal units of analysis are called Elementary

Discourse Units (EDUs), which roughly correspond to clauses. For example, in Figure 1, units 1, 2 and 3 are all EDUs. These are linked to other units or larger spans by relations from a predefined inventory, such as CAUSE, BACKGROUND, CONCESSION, etc. The relations are directed and assume a distinction between more prominent “nucleus” units and less prominent “satellite” units. For example, in Figure 1, EDU 1 is satellite to the nucleus EDU 2 in the relation CONCESSION and EDU 3 is satellite to the nucleus span 1-2 in the relation CAUSE. A single hierarchical RST tree structure like that in Figure 1 is assumed to hold over an entire document. RST was used to annotate the RST Discourse Treebank (RST-DT), which contains a large portion of the same Wall Street Journal data as PDTB (Carlson et al., 2003).

While RST-DT does not mark any form of signals identifying a relation, such as the presence of connectives, subsequent work on the RST Signalling Corpus (RST-SC, Das and Taboada 2018) proposed to address this lack of signal annotations by adding signal-type annotations to the relations in RST-DT. However, RST-SC did not anchor the signals to specific tokens.

Two key insights from PDTB and RST-SC are the importance of identifying exact tokens indicative of relations, and the possibility that signals may indicate the presence of multiple, concurrent relations between units, which would not be representable in a strict tree structure. As a result, Zeldes et al. (2025) proposed an enhanced version of RST (eRST), which anchored RST relations to token-aligned signals in 8 major categories, and allowed for the addition of multiple, concurrent and tree-breaking edges where signals indicate their presence. These relations, called ‘secondary edges’, were added to the existing RST trees in the Georgetown University Multilayer corpus (GUM, Zeldes 2017), resulting in increased coverage of relations that would be missed in basic RST due to tree constraints.

2.3 GDTB

GDTB (the Georgetown Discourse Treebank) is a PDTB-style corpus semi-automatically converted from the GUM corpus (Liu et al., 2024), following the latest PDTB guidelines. In this paper, we will leverage the existence of the GDTB annotations in order to identify a. implicit connectives corresponding to GUM eRST relations, and b. cases where additional relations signaled by such a connective

may be missing.

Since the present work is focused on adding implicit relations in GDTB to eRST, it is necessary to explain how implicit relations in GDTB were derived. PDTB guidelines allow for implicit connective in three cases: between adjacent sentences in the same paragraph; between coordinate predications within the same verb phrase; before adverbial clauses not introduced by a conjunction. GDTB annotations follow these guidelines: First, every junction allowing for an implicit relation (where a connective could be inserted) was identified using syntax trees, which are available from the GUM corpus. Second, a connective prediction model was trained on implicit relations from the PDTB training set to output a list of potential connectives at that juncture. Finally, the connectives and any RST relations applying at that juncture were fed to a pretrained relation classifier system (DisCoDisCo, Gessler et al. 2021) to select the most likely PDTB relation. The complete test set of the corpus and a select set of relation types determined to be unreliable in the entire corpus were then manually corrected. Overall, the system reaches an F1 score of 0.8485 on exact label and span match and an F1 score of 0.9489 on span match only when tested on the test set. Below, we leverage these annotations and provide a comprehensive evaluation of our system in section 5.

2.4 Mapping PDTB to RST-DT

Demberg et al. (2019) proposed the first alignment algorithm for aligning PDTB to RST. This alignment faces two main challenges. First, the two frameworks differ in segmentation of argument spans. Second, there is no one-to-one correspondence between PDTB labels and RST labels. To address the segmentation mismatch, they drew on RST’s Strong Nuclearity hypothesis (Marcu, 2000), which states that when a relation is postulated to hold between two spans of text, it should also hold between the nuclei of those spans. Accordingly, even when the two arguments of a PDTB relation do not correspond exactly to the two spans involved in a given RST relation, the relations can still be aligned if the PDTB spans match the head nuclei of the two RST constituents. In following this, they were able to map 76% of PDTB relations to RST relations, out of which 52% percent had perfect alignment and the remaining 48% relied on the Strong Nuclearity hypothesis. They further found that agreement with theoretically proposed

PDTB–RST label correspondences was low, at less than 50%, and provide empirically attested mappings for implicit relations. Concurrently with that work, Bourgonje and Zolotareno (2019) attempted to induce PDTB relations from RST trees semi-automatically in German. Later, Costa et al. (2023) aligned both explicit and implicit relations in English PDTB 3.0 to RST-DT and noted that implicit relations pose a greater challenge for sense mapping than explicit relations. Mapping PDTB-style annotations to RST-style ones is also an important part of the present work, which faces similar challenges.

3 Data

The present study seeks to build on the GUM corpus Version 12 (Zeldes, 2017), which includes 16 formal and informal, spoken and written English text types. It originally contained RST style annotations but has been enhanced using eRST. With over 290K tokens, this is currently the largest RST dataset concurrently annotated for DMs and non-DM signals of discourse relations.

All of our implicit connectives for the data come from the implicit relations in GDTB. Liu et al. (2024) evaluate the implicit relations in the corpus and establish an 0.94 F1 score for correct argument spans. Although on the task of exact PDTB labeling (span, relation type, direction and connective) the corpus reaches a lower score of 0.85, for our purposes, most of the tasks besides predicting the direction of added relations do not require exact PDTB label predictions. Where possible, the GDTB corpus also provides the original RST relation that is converted to the PDTB relation, unless no such relation exists, further facilitating our task.

4 Adding implicit connectives

Figure 2 illustrates what we add to the original eRST-style annotations. GDTB suggests an implicit connective **AND** between units 2-4 and 5. It is added to EDU 5 and anchored to the relation **JOINT-LIST** above 5 as an implicit signal because “and” is compatible with that relation. Details about the anchoring algorithm are provided in Section 4.1 below. On the other hand, GDTB suggests an implicit connective **AS A RESULT** between units 1 and 2-5. However, it cannot be anchored to the relation **EXPLANATION-EVIDENCE** in the original RST tree because the connective “as a result” suggests that the sentences in 2-5 are results of 1, whereas in

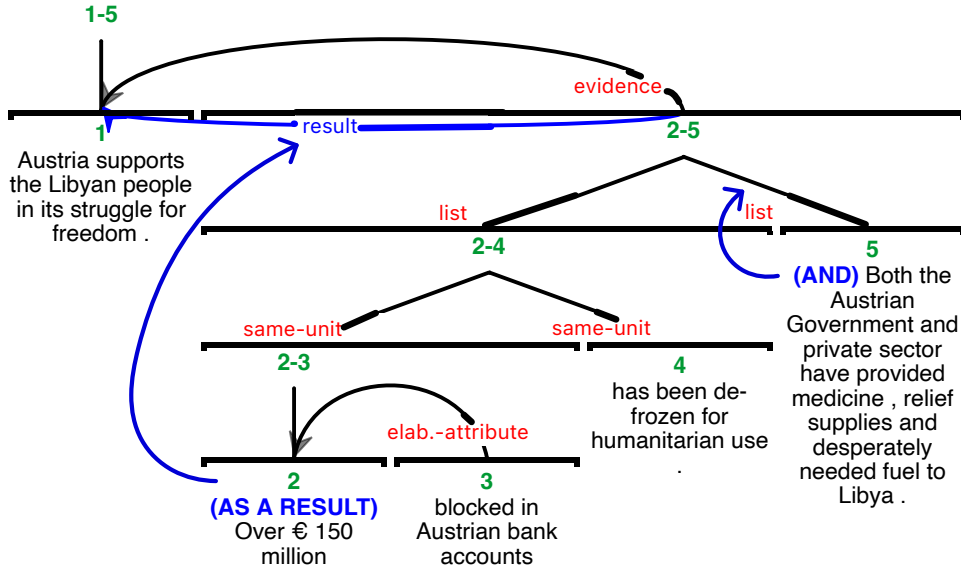


Figure 2: An eRST tree with implicit connectives and an added secondary edge in blue. Example adapted from GUM_speech_austria, EDUs 19-23.

the primary tree, considering the overall structure of the text, 2-5 are evidence for Austria’s support. Thus, a new secondary edge CAUSAL-RESULT is added from 2-5 to 1, to which “as a result” is anchored as an implicit signal. The process of adding secondary edges is detailed in section 4.2 below.

4.1 Anchoring to the Primary RST Tree

Our algorithm takes the GDTB data as a starting point because by following PDTB v3 guidelines, we expect to find all relations corresponding to the PDTB definition of implicit connective-based relations in its data. We start from the argument spans associated with the implicit connective in GDTB, and try to map them to the original RST listed in GDTB as its parallel, if one is available. The relation label between the chosen RST segments should match the corresponding RST relation between the GDTB arguments based on the mapping proposed by Liu et al. (2024).

4.1.1 Perfect Match

The most basic case is where there is a perfect token-span correspondence between GDTB arguments and RST constituents. eRST relations available for implicit signal anchoring include satellite-nucleus relations, multinuclear relations and source-target secondary edges, which are directed but do not express nuclearity. When the GDTB arguments parallel an RST multinuclear relation (when they share a multinuclear parent node

as in unit 2-5 in Figure 2), we convert the multinuclear relation into a dependency relation from the argument on the right to the argument on the left, following the dependency conversion by Li et al. (2014). For example, the list relation in the figure could correspond to a PDTB conjunction relation spanning the tokens in 2-4 as Arg1 and unit 5 as Arg2. This facilitates the mapping between GDTB and eRST.

4.1.2 RST spans contain GDTB arguments

GDTB	Text	RST
1	I grew up in a... public service family.	N
2	(because) My father... was mayor of Baltimore.	N
	My brother was mayor later... in Baltimore.	N

Figure 3: GDTB arguments contained within RST: an example from GUM_speech_floyd, EDUs 59–61. Connective in red is suggested by GDTB and added to RST. The shadowed part illustrates the RST relation the connective is anchored to.

Due to PDTB’s minimality principle, which states that only the minimal text needed to interpret a discourse relation should be included in the argument spans (Prasad et al., 2014), the GDTB authors clipped extra materials from RST spans to create GDTB arguments. See Figure 3 for an example: GDTB identifies row 1 and row 2 as constituting the relation triggered by an implicit

because. However, RST suggests that row 2 first forms a JOINT-LIST relation with row three, and then an EXPLANATION-EVIDENCE relation (which corresponds to the connective “because”) with the first row as the nucleus.

As a result, the two GDTB arguments in Figure 3 do not themselves form a valid RST relation. Both correspond to nuclei. However, the parent of the EDU corresponding to GDTB Arg2 stands in the relevant RST relation to the EDU corresponding to GDTB Arg1. In such cases, where GDTB arguments correspond to subspans of larger RST segments, we recursively traverse upward from the RST spans corresponding to the GDTB arguments and inspect the relation at the parent level.

We begin by traversing upward from the left span. At each step, we test whether the relation between the current parent of the left span and the right span matches an RST compatible with the connective. If not, we continue ascending to the next parent, stopping once we hit the root. We then repeat the same search using the parent of the right span, again traversing through the ancestor constituents of the left span. Finally, among all candidate pairs that satisfy a compatible RST relation, we select the pair spanning the fewest EDUs.

4.1.3 GDTB arguments contain RST segments

GDTB	Text	RST
1	Find a quarter or another small coin.	N
2	(then) Have both of the teams turn around	N
	and have the referee...	
	while the teams...	

Figure 4: RST contained in GDTB: an example from GUM_who_wquidditch, EDUs 128-131

On the other hand, sometimes GDTB arguments include extra material not contained in the RST relation they correspond to because PDTB arguments of intersentential relations tends to cover full sentences. For example, in Figure 4, GDTB includes three rows as the second argument, which forms an implicit relation with the first row triggered by the implicit **then**. In RST, on the other hand, the first row and only the second row forms the JOINT-SEQUENCE relation, which corresponds to the “then” connective.

In cases where the GDTB relation includes extra material outside the RST segments, we trim down the GDTB arguments to see if any sub-parts can

align with the RST relation to which Liu et al. map the GDTB relation. Because an implicit relation can only hold between two consecutive segments, the juncture between the two GDTB arguments should match the juncture between the two RST spans. Thus the logic is to preserve the core parts along the juncture and trim down the end points.

The trimming process begins at the left boundary of the left argument, while the right argument is held fixed; starting instead from the right boundary would be equivalent. At each step, we remove the leftmost EDU and test whether the desired RST relation holds between the remaining portion of the left argument and the right argument. We continue this process until there are no more EDUs in the left argument. We then remove the rightmost EDU from the right argument and restart the same procedure on the left argument, again trimming from the left boundary. For each progressively shortened right span, we trim the left span from the left. The process terminates once a matching relation is found, if one exists.

4.1.4 Overlapping RST and GDTB

GDTB	Text	RST			
1	... the conference has got well under way,				N
	and everybody is keeping busy.				N
2	(for instance) Lionel and Bert and I all have lots to do	N			
	He is specializing more or less on Trieste,		N	S	
	while Bert and I cover pretty much everything else.			N	S

Figure 5: Partial GDTB contained in RST: an example from GUM_letter_conference, EDUs 5-9

Finally, there are also cases where only subparts of the GDTB arguments correspond to a valid RST relation, which also includes extra material outside of the GDTB arguments. For example, in Figure 5, **for instance** is thought to trigger an implicit relation between row 1–2 and row 3. However, the corresponding RST relation ELABORATION-ADDITIONAL holds between row 2 and row 3–4. Thus, there is some overlap between the GDTB arguments and the RST arguments, and compatible label+connective, but no containment.

For such cases, we use the simple heuristic to pick out the largest adjacent RST segments contained in the GDTB arguments as a base case, and apply the recursive algorithm of 4.1.2 to determine the appropriate RST spans.

GDTB	Text	RST
1	that	N
2	(in order) to understand how fear works in us,	S
1	we have to see it in the context	N
	where it evolved.	

Figure 6: Same-unit: an example from GUM_essay_evolved, EDUs 17-20

4.1.5 Trouble cases: Same-unit

Intra-sentential implicit relations pose a challenge to our algorithm so far because one GDTB argument can be contained within the other, as shown in Figure 6 and also the desired RST relation does not necessarily hold between the arguments along the juncture.

In such cases, one GDTB argument is contained within another because discontinuous EDUs can be connected via the RST pseudo-relation SAME-UNIT also shown in Figure 2. We then check the RST relation between the RST span that the contained segment corresponds to, and that of either outer layer to see if it is compatible with the implicit connective. For example, in Figure 6, GDTB argument 2 is contained between the two GDTB argument 1s. We would thus check if there is a valid RST relation between argument 2 and either one of argument 1s. In this case, we find that the relation PURPOSE-GOAL holds between the satellite (argument 2 in row 2) and nucleus (argument 1 in row 3–4), and matches the connective *in order*.

4.2 When and How to Add New Edges?

After we locate the correct RST relation, we need to decide whether the implicit connective can serve as a signal for that relation. In GDTB, the connective is predicted given two spans of text, so it may represent a more local relation incompatible with the RST relation it is mapped from, as shown in the connective “as a result” in Figure 2.

We decide on whether an implicit connective is a signal for a particular RST relation based on whether it is ever aligned with that relation as an explicit DM signal in the eRST corpus. If the connective is unattested for that particular RST relation, a secondary edge is proposed between the two RST segments, indicating that the relation signaled by the implicit connective is of a different kind than the existing RST relation.

In order to add the correct secondary edge, we

use a combination of rule-based mapping and predictions from an LLM, using the fine-tuned Qwen3-4B decoder model DeDisCo (Ju et al., 2025), a discourse relation classifier which remains state-of-the-art on the DISRPT shared task benchmark for relation classification (Braud et al., 2024, 2025).

Since DeDisCo performs best when given the direction of the discourse relation, we first get the RST direction from the GDTB sense. Although there is not a one-to-one mapping between the RST direction and GDTB arguments because what GDTB calls Arg1 and Arg2 depends on the syntactic configuration and their linear order in the text, for some cases the RST direction can be reliably mapped from the GDTB sense. Since GDTB uses PDTB v3 labels, Level 3 senses indicate which argument of the two plays what role in the relation, giving clues as to which argument is likely to be the satellite or nucleus in RST. For example, what an RST CONCESSION relation concedes is reliably the opposite argument span of the .ARGX-AS-DENIER argument in GDTB. For other labels, we stipulate the direction to be the most common one in the corpus for the relation and connective combination. For example, all the senses that have level 1 as temporal have the RST relation from Arg2 to Arg1 if they are signaled by “then”. We then specify the corresponding direction in the prompt and use DeDisCo to predict the discourse relation. In the 2025 DISRPT shared task, fine-grained RST relations are collapsed into 17 coarse-grained discourse labels, so the label DeDisCo predicts may correspond to multiple possible RST relations.

We get the probabilities from the DeDisCo predictions on the DISRPT labels and first take the label with the highest probability, then map it to possible RST relations. Among the possible RST relations, we again resort to the mapping between explicit connectives and RST relations and choose the one most often signaled by our connective and add it as a secondary edge between the RST spans we found in section 4.1, following the direction we mapped above. If it is not possible to add any of the relations because none of them is signaled by our connective, we move on to the label with the second highest probability and repeat the same process.

After we add the optimal secondary edge, we then anchor the implicit connective to it as a signal.

5 Evaluation

To assess the quality of our corpus, we first conduct an inter-annotator agreement study between the authors, who independently edited the system outputs on 16 of the 32 test documents. Second, we evaluate the algorithm itself against the manual corrections produced by one author on the full set of 32 test documents. Table 1 gives an overview of the results.

Setting	Task	P	R	F1
IAA	Insert connective	0.99	0.99	0.99
	Connective choice (μ)	0.96	0.96	0.96
	Secedge location	0.91	0.89	0.90
	Secedge relation (μ)	0.96	0.96	0.96
	Anchoring	0.97	0.97	0.97
Algo Eval	Secedge location	0.85	0.88	0.86
	Secedge relation (μ)	0.96	0.96	0.96
	Anchoring	0.99	0.99	0.99

Table 1: IAA and model performance. Micro-averaged scores (μ) are used when label classes are imbalanced.

5.1 Inter-annotator agreement

We measure IAA for three subproblems. First, we assess whether annotators agree that an implicit connective can felicitously be inserted at the position suggested by GDTB. Restricting attention to cases where both annotators agree that a connective should be inserted, we measure the agreement rate on connective choice. Second, we evaluate the IAA on the addition of new secondary edges triggered by the implicit connective, including both span selection and edge direction. For cases where both annotators agree that a secondary edge should be added, we measure the IAA on relation labeling. Third, we evaluate whether implicit connectives are correctly anchored to the original eRST relation and to the newly added secondary edge in cases where both annotators agree on the added edge. Since this annotation task corresponds to a detection scenario (the number of annotation decisions is not pre-determined), we use mutual precision, recall and F1 scores between the annotators as the metric, rather than kappa (P and R are interchangeable, depending on which annotator is taken as the reference). Overall, the F1 scores indicate that the annotation set-up is highly reliable.

5.2 Algorithm Evaluation

Since our method directly adopts implicit connectives from GDTB, the algorithm is not responsible

for predicting whether a connective should be inserted or which connective should be used. We therefore evaluate the algorithm only on the second and third task mentioned in Section 5.1. Overall, the algorithm shows strong performance, with particularly strong results on connective anchoring once the additional secondary edges are added. The comparatively lower score on secondary edge location is due to the difficulty in assigning the right direction for a secondary edge based on the directionless GDTB annotations.

6 What role do implicit connectives play?

In this section we will answer several questions about the distribution and possible functions of implicit connectives in identifying discourse relations within the framework of eRST.

How often do implicit connectives mark relations? Implicit connectives are aligned to 5320 primary relations, or 15% of the total primary relations in the corpus, and 1045 secondary edges in total, including the ones that we added in this work (42% of the total secondary edges).

How often do implicit connectives identify additional relations? In 1003 cases, implicit connectives could not be aligned to an existing edge in the data, resulting in the introduction of secondary edges. This represents an increase of 3% in the total number of edges in the corpus and an increase of 69% in the total number of secondary edges, suggesting that PDTB implicit connectives trigger many additional tree-breaking relations in eRST.

Which connectives are most common? Figure 7 shows, for each category, the six most frequent implicit connectives, with their percentages displayed side by side. Table 2 in Appendix B reports the corresponding raw counts for these connectives in each category. Connectives like “then” and “in order” are predominantly anchored to existing relation, likely because they are easily triggered by other cues: “then” can be inserted whenever there is a temporal “and”, and “in order” whenever there is a purpose infinitival “to”. On the other hand, “in fact” is more commonly used for added secondary edges, often between sentences. This may be because “in fact” can signal either contradiction or elaboration with respect to the previous context, making it particularly suitable for cases related to CONTEXT-CIRCUMSTANCE, an RST relation without a close PDTB counterpart and therefore unable

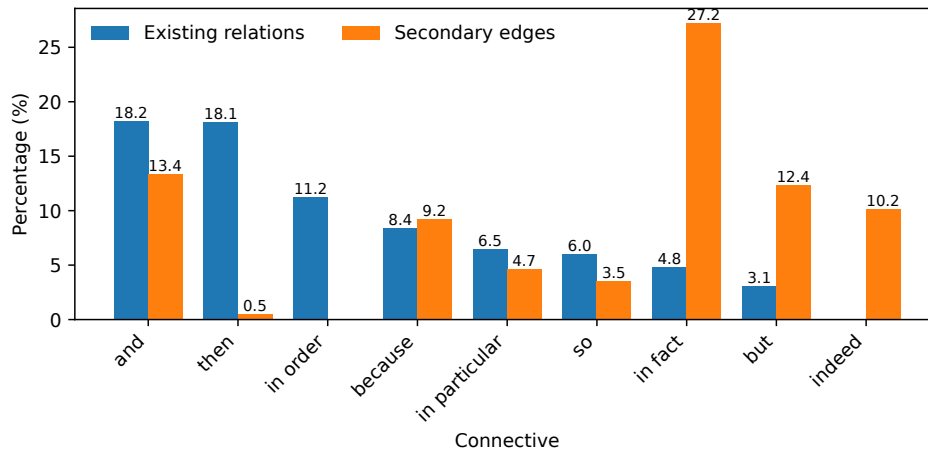


Figure 7: Most common implicit connectives anchored to existing relations vs. added secondary edges.




Figure 8: Pearson residuals for the contingency table of signal type by category (implicit vs. explicit). Circle size and color indicate the magnitude and direction of deviation from independence. *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$

to be associated with an implicit connective.

What signals coincide with implicit connectives?

Previous research has not established which types of signals are more or less likely to co-occur with implicit connectives. Here we perform an elementary analysis on the distribution of co-occurring signal types. Such investigation can help clarify which cues support discourse-relation interpretation in the absence of explicit discourse markers. We note that, once added to the graph, edges added because of implicit connectives in this work are also eligible for carrying additional non-DM signals according to eRST. To add these, we use the eRST signaling toolkit described in (Zeldes et al., 2025), which is reported to be 87.9% accurate, resulting in additional signals for implicit secondary edges in 30% of cases.

There are 6,365 relations signaled by an implicit connective and 8,300 by an explicit discourse marker. Implicit connectives co-occur with 3,627 non-dm signals (57%), whereas explicit dms co-occur with 2,192 non-dm signals (26%). It

makes sense that implicit connectives co-occur with more non-dm signals on average, since they are not overtly realized in the text and thus more co-occurring signals may be required to establish their discourse function.

We tested whether signal-type distributions differ between implicit connectives and explicit connectives (dms) using a chi-square test of independence over seven signal categories. For a detailed illustration of the seven signal types, see Appendix A. The association is significant, $\chi^2(6, N = 5819) = 301.46, p < .001, V = 0.242$. Figure 8 shows that implicit connectives are more likely to co-occur with reference signals and much less likely with morphological signals. Reference signals, including comparatives, demonstratives, pronouns, and propositional anaphora, likely prompt annotators to establish a relation between entities or propositions across adjacent segments, thereby motivating the insertion of a discourse marker. On the other hand, morphological signals, such as use of tense sequence to indicate temporality, might not moti-

vate the addition of an implicit connective.

7 Discussion

We introduce PDTB-style implicit connectives as an additional type of eRST signal. These connectives may support an existing RST relation by making its interpretation more explicit. For example, although both the CONTRAST relation in (3) and the RESULT relation in (4) are anchored to “instead”, the relations themselves are more precisely captured by the implicit connectives “but” and “so”, respectively. They may also trigger secondary edges, which capture more local semantic relations, whereas primary RST relations are organized with respect to the overall communicative intent of the author. For example, in Figure 2, “as a result” captures the relation CAUSAL-RESULT that may exist between the two sentences independently of the overall discourse tree. As a result, eRST integrates more local, tree-breaking discourse relations into the RST tree formalism, yielding a more comprehensive account of the discourse relations in a text. Finally, introducing implicit connectives as a distinct type of eRST signal allows us to examine the interaction between connectives (a.k.a. discourse markers) and non-DM signals more comprehensively, beyond the case of overt signals alone.

References

- Peter Bourgonje and Olha Zolotarevko. 2019. [Toward cross-theory discourse relation annotation](#). In *Proceedings of the Workshop on Discourse Relation Parsing and Treebanking 2019*, pages 7–11, Minneapolis, MN. 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 1–20, Suzhou, China. 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.
- Lynn Carlson, Daniel Marcu, and Mary Ellen Okurowski. 2003. Building a discourse-tagged corpus in the framework of Rhetorical Structure Theory. In *Current and New Directions in Discourse and Dialogue*, Text, Speech and Language Technology 22, pages 85–112. Kluwer, Dordrecht.
- Nelson Filipe Costa, Nadia Sheikh, and Leila Kosseim. 2023. Mapping explicit and implicit discourse relations between the RST-DT and the PDTB 3.0. In *Proceedings of the 14th International Conference on Recent Advances in Natural Language Processing*, pages 344–352.
- Debopam Das and Maite Taboada. 2018. RST signalling corpus: A corpus of signals of coherence relations. *Language Resources and Evaluation*, 52(1):149–184.
- Vera Demberg, Merel CJ Scholman, and Fatemeh Torabi Asr. 2019. How compatible are our discourse annotation frameworks? insights from mapping RST-DT and PDTB annotations. *Dialogue & Discourse*, 10(1):87–135.
- 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.
- Luke Gessler, Yang Liu, and Amir Zeldes. 2019. [A discourse signal annotation system for RST trees](#). In *Proceedings of the Workshop on Discourse Relation Parsing and Treebanking 2019*, pages 56–61, Minneapolis, MN. Association for Computational Linguistics.
- Zhuoxuan Ju, Jingni Wu, Abhishek Purushothama, and Amir Zeldes. 2025. [DeDisCo at the DISRPT 2025 shared task: A system for discourse relation classification](#). In *Proceedings of the 4th Shared Task on Discourse Relation Parsing and Treebanking (DISRPT 2025)*, pages 48–62.
- Sujian Li, Liang Wang, Ziqiang Cao, and Wenjie Li. 2014. [Text-level discourse dependency parsing](#). In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 25–35, Baltimore, Maryland. 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*, 8(3):243–281.

Daniel Marcu. 2000. *The theory and practice of discourse parsing and summarization*. MIT press.

Mitchell P. Marcus, Beatrice Santorini, and Mary Ann Marcinkiewicz. 1993. [Building a large annotated corpus of English: The Penn Treebank](#). *Computational Linguistics*, 19(2):313–330.

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.

Manfred Stede. 2012. *Discourse Processing*. Synthesis Lectures on Human Language Technologies 4. Morgan & Claypool, [San Rafael, CA].

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.

A eRST Signals

B Raw Counts of Implicit Connectives Aligned with Existing eRST Relations vs. Secondary Edges

Existing Relations	Count	Added Secedges	Count
<i>then</i>	776	<i>in fact</i>	175
<i>and</i>	726	<i>and</i>	147
<i>in order</i>	526	<i>but</i>	80
<i>because</i>	313	<i>because</i>	74
<i>in particular</i>	251	<i>indeed</i>	64
<i>so</i>	239	<i>in particular</i>	33

Table 2: Most common implicit connectives aligned with existing eRST relations and secondary edges added in this paper.

signal type	subtypes	example
dm	but, then, on the other hand...	[They wanted to] [but couldn't] <adversative-contrast>
graphical	colon, dash, semicolon layout items in sequence parentheses, quotation marks question mark	[Let me tell you a story :] <organization-preparation> [Introduction] <organization-heading> I wash [2 cut] <joint-list> it rained [(and snowed a bit)] <elaboration-additional> [Did you?] <topic-question> No
lexical	alternate expression indicative word phrase	e agreed [That is he said yes] <restatement-repetition> They planned a party [That s nice/Can't wait] <evaluation-comment>
morphological	mood tense	Go with them [I think you should] <explanation-motivation> I started an hour ago, [now I'm resting] <joint-sequence>
numerical	same count	[Two reasons] <organization-preparation> irst
reference	comparative demonstrative personal propositional	[I don't want it] <adversative-antithesis> I want another one They met Kim [This person / she was] <elaboration-additional> They met Kim [This encounter was] <elaboration-additional>
semantic	antonymy attribution source lexical chain meronymy negation repetition synonymy	eer is cheap , [wine is expensive] <adversative-contrast> [Kim said] <attribution-positive> they would it was funny [so they laughed] <causal-result> The house was big, [the door two meters tall] <elaboration-additional> im danced, [un didn't dance] <adversative-contrast> They met Dr. Kim [Dr. Kim/The surgeon was] <elaboration-additional>
syntactic	infinitival relative clause interrupted matrix clause modified head nominal modifier parallel syntactic construction past present participial clause reported speech subject auxiliary inversion	a plan [to win] <purpose-attribute> [I meant] <organization-phatic> I mean, a plan [to win] <purpose-attribute> articles [explaining chess] <elaboration-attribute> it's all tasty [it's all pretty] <joint-list> im appeared [dressed in black] <elaboration-attribute> [im said] <attribution-positive> that they would I would have [had I known] <contingency-condition>

Figure 9: eRST signals