A Cross-Genre Analysis of Discourse Relation Signaling in the GUM Corpus

Lauren Levine
Georgetown University
lel76@georgetown.edu

Abstract

In this paper, we investigate the cross-genre variation in how discourse relations are signaled in the Georgetown University Mutilayer (GUM) Corpus, an English language corpus which contains 16 different genres of texts with various linguistic annotations, including Rhetorical Structure Theory (RST) style discourse annotations. We look at the proportions of discourse signals in each genre, and then we conduct an analysis of which discourse relations display the most inter-genre variation in how they are signaled, providing a methodology for ranking the inter-genre variability of the signaling of individual discourse relations. Although the way in which individual discourse relations are signaled in GUM is relatively stable across genres, we are able still to produce stable rankings, finding that `organization`, `restatement`, and `explanation` relations display the most inter-genre variation. However, we find that genre specific graphical norms can account for a large portion of the observed variation.

1 Introduction

Discourse relations are used to describe the meaning that arises from the combination of multiple linguistic units in a discourse. In computational discourse analysis, there have been multiple linguistic formalisms proposed regarding how to annotate this phenomenon, each with their own unique inventories of discourse relations. One such prominent formalism is Rhetorical Structure Theory (RST; Mann and Thompson, 1988), which assigns relation labels on a pragmatic basis, without reference to particular linguistic signals. Despite this, previous work on signaling in RST data has found that over 90% of RST discourse relations are signaled in some way (Das and Taboada, 2018b). This includes overt discourse markers, as shown in Figure 1 with the explicit discourse marker *because*, as well as other more implicit discourse signals. As such, we many wonder if there are patterns in the distributions of which types of discourse signals appear with different discourse relations, and if so, how these patterns appear across different genres. Answers to such questions will provide insight into whether a pragmatic formalism like RST also displays structural patterns in its annotation, which would not necessarily be expected as there are not sturctural criteria in the annotation of RST discourse relations. Such investigation will also provide insights for describing genre differences, particularly regarding how genres use different structural means to achieve a particular discourse purpose.

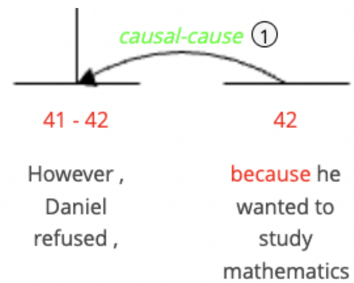

Figure 1: Example from the GUM corpus of a `causal` discourse relation, overtly signaled by the explicit discourse marker (dm) *because*.

In this paper, we see that different RST relations do in fact co-occur with different proportions of discourse signal types (Figure 2), and we focus in on the question of how different genres signal the same discourse relations. We then consider which individual RST relations display the most inter-genre variation in how they are signaled. In order to investigate this, we introduce an inter-genre variation ranking metric: average pairwise Jensen-Shannon distance (Avg. Pairwise JSD), the details of which are given in Section 4. After we use this metric to obtain an inter-genre variation ranking for different discourse relations, we explore the rela-

Proceedings of the Society for Computation in Linguistics (SCiL) 2025, pages 300-308.
Eugene, Oregon, July 18-20, 2025

tions with the most variability to see what this can tell us about the characteristics of different genres.

Overall, we find that while there is clearly inter-relation variation in the distributions of discourse signals, the means of signaling individual RST relations remain relatively consistent across genres. This indicates a general stability in the manner of signaling individual discourse relations, which we would not necessarily expect considering that RST is a pragmatic formalism. However, by utilizing the Avg. Pairwise JSD metric, we are still able to produce stable rankings for which discourse relations show the most cross-genre variation how they are signaled, finding that `organization`, `restatement`, and `explanation` relations display the most inter-genre variation, and that `evaluation` and `adversative` relations show the least inter-genre variation. Code and visualizations for this paper are available on GitHub[1].

2 Previous Work

While relation signaling in the RST formalism is a relatively new area of interest, there are several foundational works which we draw upon in this investigation. Firstly, a major resource for RST data in English is the RST Discourse Treebank (RST-DT), which consists of 385 Wall Street Journal articles (Carlson et al., 2002). In 2013, Taboada and Das subsequently added an additional layer of signaling annotations to a portion of this corpus, and later the entire RST-DT corpus, creating the RST Signalling Corpus (RST-SC; Das and Taboada, 2018a). This work provided the first available RST data with signaling information, and established a manageable taxonomy of signal types, including not only overt discourse markers, but various implicit discourse signals as well. Since its creation, the RST-SC has been used for various corpus analyses of relation signaling (Das and Taboada, 2018b; Das, 2019; Egg and Das, 2022).

There have also been a number of efforts aimed at extending the application of the relation signaling framework created by Taboada and Das. As RST-SC does not indicate which tokens are aligned with the signal type annotations, Liu and Zeldes (2019) made efforts to anchor signaling information directly to tokens in a text. Additionally, Gessler et al. (2019) created an online annotation tool for adding signaling information directly onto existing RST annotations. Both of these efforts were further leveraged in the creation of signaling annotations in data for Enhanced Rhetorical Structure Theory (eRST), an extension of the theoretical RST framework which added a means to account for "tree-breaking, nonprojective and concurrent relations" in discourse relation graphs (Zeldes et al., 2024). The eRST project follows the relation signaling taxonomy from Taboada and Das, dividing relation signals into the following categories: discourse markers, graphical, lexical, morphological, numerical, reference, semantic, and syntactic. The corpus analysis we conduct in this paper is focused on the signaling annotations added to the GUM RST treebank from the eRST project.

3 Data

For this investigation, we use GUM Version 10[2], a 228k token corpus of English, which is composed of 235 documents, divided approximately evenly across 16 different genres: academic, biographies, courtroom, conversation, essay, fiction, interview, letters, news, podcasts, speeches, textbooks, travel, vlogs, how-to and Reddit forum discussions (Zeldes, 2017). As mentioned in the previous section, the GUM corpus has signaling annotations consistent with the form established for the eRST formalism, extended from the taxonomy created by Taboada and Das.

For this analysis, we only consider discourse relations which co-occur with at least one signal annotation (at all levels of the eRST tree). There are a total of 30,774 discourse relation annotations in GUM v10, 69.35% of which (21,343 instances) occur with one or more signaling annotation. The eRST annotations in GUM leverage a two-tiered relation inventory, where the coarse relation and the fine-grained subtype are connected with "-" (e.g., `causal` is the coarse relation type for the fine-grained relation `causal-cause`). The full relation inventory of 15 coarse relations and 32 fine-grained relations is shown in Appendix A. For each relation signal annotation, we extract the signal type and signal subtype, the RST relation type and RST relation subtype (e.g., `elaboration` and `elaboration-attribute`), and the genre in which it occurs from the GUM corpus. This means that a single relation will be extracted multiple times if it occurs with multiple signals. And while we extract both the signal type and the signal subtype, in order

[1] https://github.com/lauren-lizzy-levine/gum-genre-signaling

[2] https://github.com/amir-zeldes/gum

to have enough instances in each signal category to analyze statistically, we limit our investigation to the higher level signal types: discourse markers (dm), graphical (grf), lexical (lex), morphological (mrf), numerical (num), reference (ref), semantic (sem), and syntactic (syn). For reference, the complete signal inventory from Zeldes et al. (2024) is included in Appendix A. For RST discourse relations, we investigate at the level of both coarse relations (e.g., `elaboration`) and fine-grained relations (e.g., `elaboration-attribute`) from the RST relation inventory.

4 Methods

In order to investigate the inter-genre variability of signaling for individual relations, we need a means of quantifying how different the distributions of relation signals are between a pair of genres for a particular relation. We adopt the Jensen-Shannon Distance[3] as metric for this purpose.

The Jensen-Shannon Divergence (JS-Div) is a symmetric measure of the similarity of two probability distributions. This metric is bounded, $0 \leq$ JS-Div ≤ 1, where 0 indicates the distributions are identical, and 1 indicates they are completely different. The Jensen-Shannon Distance (JSD) is the square root of the JS-Div, and it is commonly used to assess the similarity of probability distributions. In order to apply JSD as a metric to our relation signaling data, we make the assumption that the frequency counts of the signal types used to indicate a relation in a specific genre can be used to approximate the probability distribution of how that relation is signaled in that genre[4].

For each relation (e.g., `explanation`), this gives us per genre probability distributions of signal types which we can compare using JSD. We can then calculate the JSD scores between all possible pairs of genres (e.g., ('reddit', 'academic'): 0.63, ('interview', 'academic'): 0.59, etc.). We use these scores for two purposes: First, we construct a distance matrix for genre pairs which can be used as input for clustering/dendrograms of genre similarity (with respect to signaling). Secondly, we can take an average of these scores to create a single number that represents the inter-genre variability

[3]https://docs.scipy.org/doc/scipy/reference/generated/scipy.spatial.distance.jensenshannon.html

[4]This assumption means that you take the frequency counts to be a representative distribution of the categories of signal types. This may not be a valid assumption if the data is sparse.

Rank Correlation Metric	**Relation Type**	
	Coarse	**Fine-grained**
Avg. Kendall's Tau	0.82	0.76
Avg. Spearman Rank	0.93	0.90
Avg. Pearson Correlation	0.95	0.92

Table 1: Averages of correlation metrics from comparing rankings of inter-genre variation for the signaling of individual RST relations, computed from randomly sampled subsets of the GUM corpus.

score for the individual relation (e.g., 0.35). We refer to this metric for quantifying inter-domain variation as the average pairwise Jensen-Shannon Distance (Avg. Pairwise JSD). We note that while in this study we specifically use the metric to investigation the inter-genre variation in how individual discourse relations are signaled, it can be thought of as a more general metric. Genre, relation, and signal type may be swapped out for other categories as the context requires.

The inter-genre variability score for a discourse relation R using Avg. Pairwise JSD is defined as:

$$Avg.PairwiseJSD(R) =$$

$$\frac{\sum_{i,j \in G} JSD(SD_i(R), SD_j(R))}{\binom{|G|}{2}}$$

where G is the set of genres, JSD is the Jensen-Shannon Distance, and $SD_x(R)$ is the frequency distribution of relation signal types for relation R in genre x.

For the fine-grained relations, the frequency distribution of the relation signal types is approximated by the raw frequency counts in the data. For the coarse relations, we normalize the frequency distribution by the proportions of sub-relations composing the coarse relation. We treat each sub-relation as an independent class within the coarse relation, and we take the macro-average of the distributions for the individual classes to be frequency distribution for the coarse relation.

Once the Avg. Pairwise JSD scores are calculated, they can be sorted to give a relative ranking of inter-genre variability of signaling amongst individual relations. In order to establish how reliably the Avg. Pairwise JSD is able to construct this relative ranking, we compare the rankings that this methodology produces when applied to different subsets of the data. For each genre, we randomly sample 5 documents, and we compute the relative rankings of inter-genre variability as described

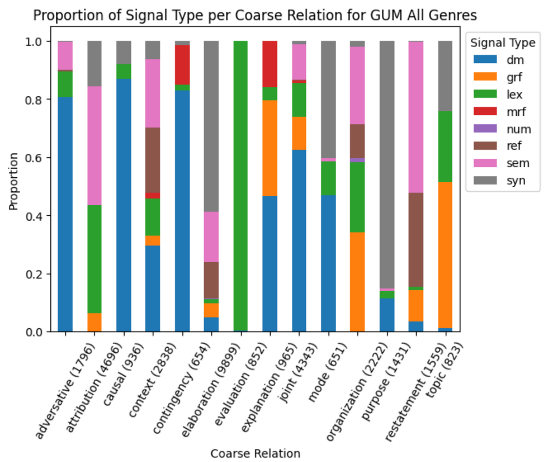

Figure 2: Proportions of relation signal types for coarse RST relations in the GUM corpus. The total number of occurrences of a relation type co-occuring with a signal is included in parentheses.

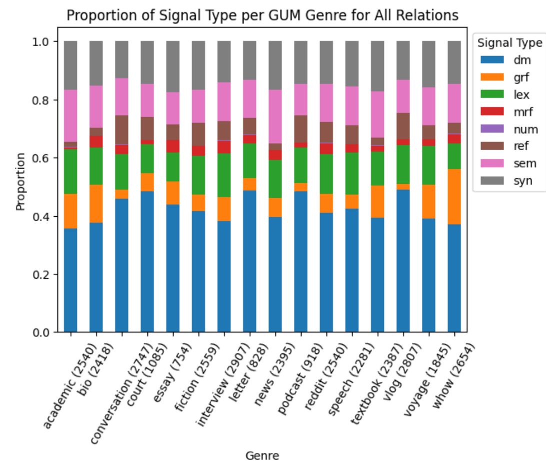

Figure 3: Proportions of relation signal types for the genres in the GUM corpus. The total number of occurrences of relations co-occuring with a signal in the given genre is included in parentheses.

above. We repeat this process 50 times, so there are 50 independent rankings (each with Avg. Pairwise JSD scores) for both coarse and fine-grained RST relations. For each pair of rankings in this 50 run sequence (1225 pairs), we calculate the following correlation metrics between the rankings: Kendall's Tau, Spearman Rank and Pearson Correlation, and then we average the resulting scores for each.

We report the averages for the rank correlation metrics Table 1. For all of these metrics, the closer the score is to 1, the closer the correspondence between the rankings/scores being compared. We see that all the metrics are quite high, and that the metrics for coarse relation ranking averages are higher than those for the fine-grained relation rankings. The strength of the correlation coefficients shows that the rankings are relatively stable, even when data is randomly sampled. We take this to be a reasonable indication that Avg. Pairwise JSD can be used to reliably construct a relative ranking of inter-genre variation for individual relations.

5 Results

To begin our analysis, we investigate the variation in signal types used for different relations in the GUM corpus. Figure 2 visualizes the proportions of signal types used with each coarse RST relation in the GUM corpus. We see that there is a considerable amount of inter-relation variation, and there are some interesting observations to be made from this visualization alone: `evaluation` relations are signaled exclusively by lexical features, `adversative`, `causal`, and `contingency` relations are dominated by overt discourse markers, etc.

However, as this investigation is focused on the inter-genre variation of individual relations, we shift our focus to explore the distribution of relations signals across the different genres in the GUM corpus. We provide a visualization for this analysis in Figure 3, which shows the proportions of signals present in each genre, adjusted for the relative frequencies of the relations present in that genre[5].

In Figure 3, we see that the signal proportions are surprisingly consistent across the various genres of the GUM corpus. We face the possibility that individual relations do not display a substantial amount of inter-genre variation overall, and, as such, we need to focus in on the areas of our data which display the most inter-genre variation for investigation. To this end, we create a relative ranking of the inter-genre variability of signaling amongst individual relations via the methods described in Section 4. The inter-genre variation ranking for the coarse RST relations is shown in Figure 4, and the inter-genre variation ranking for the fine-grained RST relations is shown in Figure 5.

Looking at the ranking for coarse rela-

[5]The signal type proportions for each fine-grained relation attested in the genre are calculated separately and each one considered a separate class. The macro-average of these classes is then taken and reported in Figure 3 as the signal distribution of the genre.

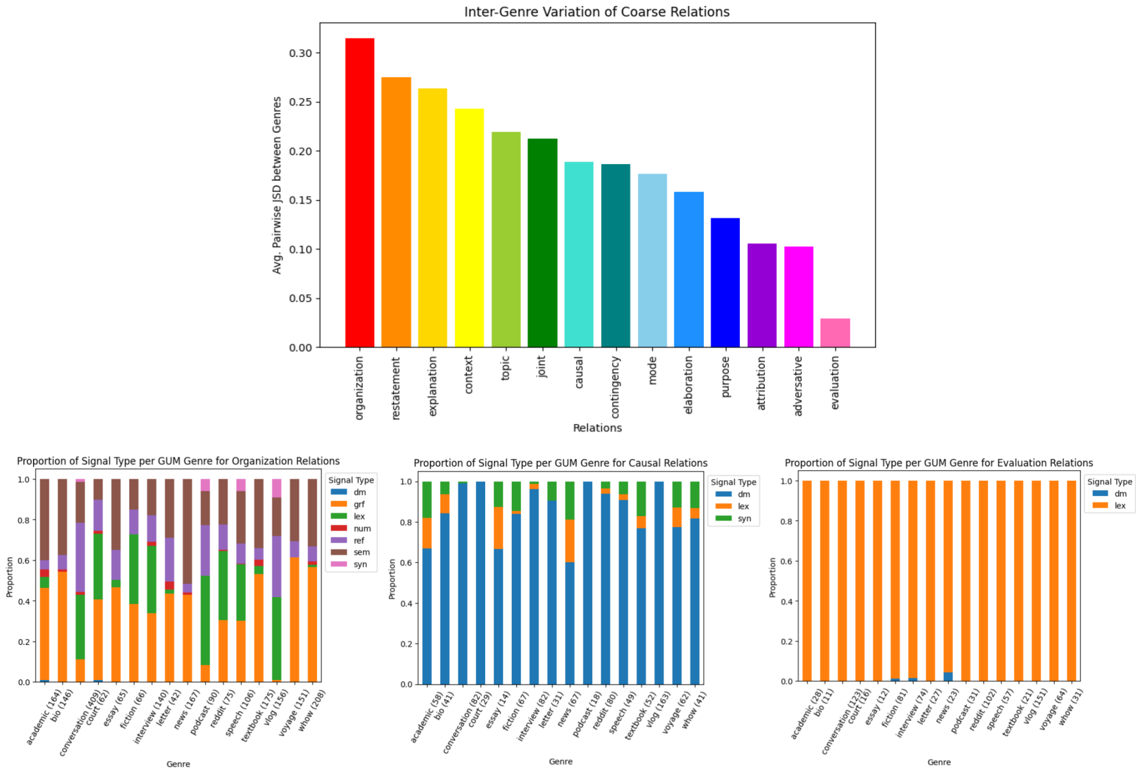

Figure 4: (Top) Ranking of inter-genre variation of relation signal distributions for coarse RST relations (based on Avg. Pairwise JSD). Proportions of relation signals across genres for: (bottom left) the coarse relation showing the most variation: `organization`, (bottom middle) the coarse relation showing the median variation: `causal`, and (bottom right) the coarse relation showing the least variation: `evaluation`.

tions in Figure 4, we see that `organization`, `restatement`, and `explanation` relations display the most inter-genre variation, while `attribution`, `adversative`, and `evaluation` relations display the least inter-genre variation. In the bottom section of Figure 4, we also show the signal type distributions across genres for the relations whose Avg. Pairwise JSD indicated that they show the most (`organization`), median (`causal`), and lowest (`evaluation`) inter-genre variation. As we can see from the visualizations, Avg. Pairwise JSD seems to accurately reflect the relative inter-genre variation of the relations.

Looking at the ranking for fine-grained relations in Figure 5, we see that `explanation-evidence`, `restatement-partial`, and `restatement-repetition` relations display the most inter-genre variation, while `elaboration-attribute`, `explanation-justify`, and `evaluation-comment` relations display the least inter-genre variation. In the bottom section of Figure 5, we again show the signal type distributions across genres for the relations whose Avg. Pairwise JSD indicated that they show the most (`explanation-evidence`), median (`adversative-antithesis`), and lowest (`evaluation-comment`) inter-genre variation. As we can again see from these visualizations, Avg. Pairwise JSD accurately reflects the relative inter-genre variation of the relations.

Now that we have rankings of the inter-genre variability for relations, we will take a look at some of the individual relations which displayed the most variation: `organization` and `explanation`. First, consider Figure 6. The left side of the figure shows the distribution of relation signal types across genres for the `organization` relation. The right side of the figure is a dendrogram showing the signaling similarity between genres for the `organization` relation (based on a distance matrix of JSD scores between genre pairs).

Looking at the dendrogram in Figure 6, we see that there is a relatively clear split between spoken genres and written genres. This means, perhaps unsurprisingly, that written genres and spoken genres are relatively distinct in how they signal

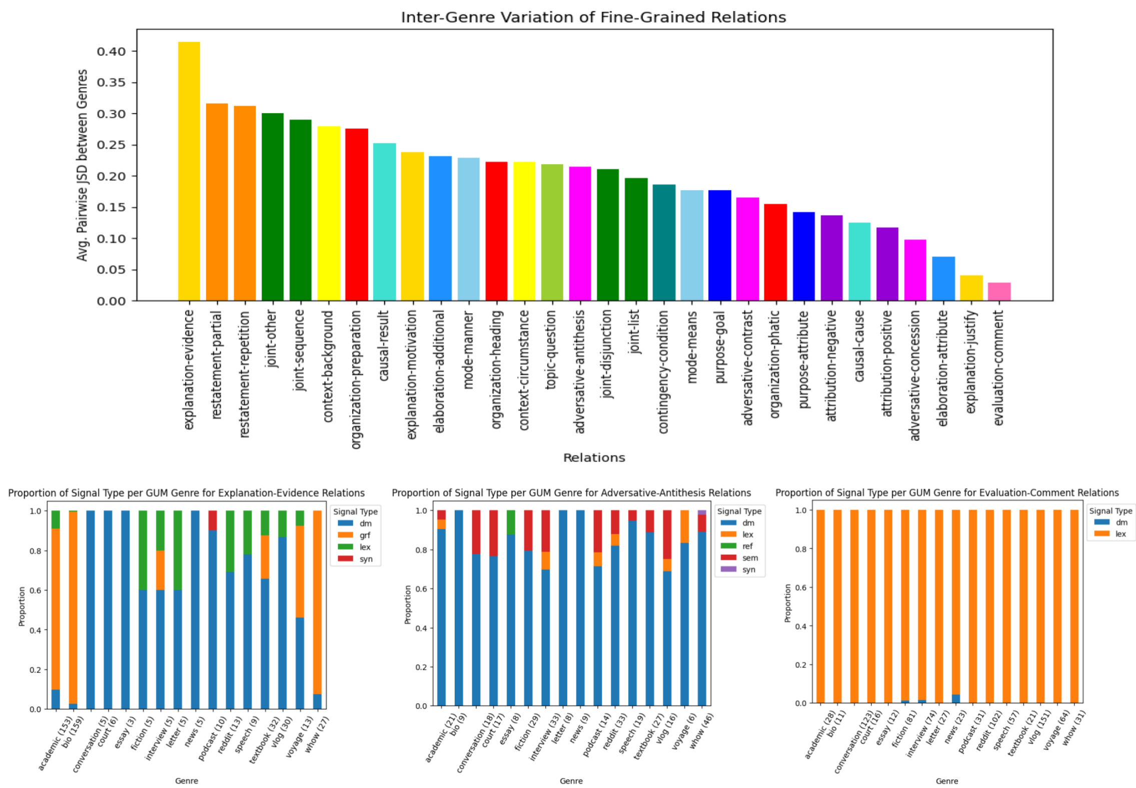

Figure 5: (Top) Ranking of inter-genre variation of relation signal distributions for fine-grained RST relations (based on Avg. Pairwise JSD). Proportions of relation signals across genres for: (bottom left) the fine-grained relation showing the most variation: `explanation-evidence`, (bottom middle) the fine-grained relation showing the median variation: `adversative-antithesis`, and (bottom right) the fine-grained relation showing the least variation: `evaluation-comment`.

`organization` relations. If we look at the graph on the left side of Figure 6, we see that spoken genres, such as conversation and podcast, have a much smaller proportion of graphical signals than the written genres. This is intuitive, as there are many graphical signals, such as headings, that are commonly used in written genres for organizational purposes, which cannot be used in spoken genres. Instead, we see that the lexical signal type compensates for the lack of graphical signals in spoken genres.

Now consider Figure 7. On the left, we have the distribution of relation signal types across genres for the `explanation` relation. On the right, we have a genre signaling similarity dendrogram, this time for the `explanation` relation. In this dendrogram, we can see that there is a clear split between academic, biographies and wiki-how, and the rest of the genres. If we look at the graph on the left of Figure 7, we once again see that different proportions of graphical signals are largely responsible for this divergence. Upon qualitative examination of the data, we see that this is largely due to parentheses being used for citations, a practice which is common in academic writing, biographies, and wiki articles.

6 Discussion

In the results of our investigation, we saw that the inter-genre signaling of individual discourse relations is relatively stable. In two of the coarse relations which showed the most inter-genre variation in their signaling, `organization` and `explanation`, genre specific graphical norms seemed to contribute more to the existing variation than the language content. As such, if there is a large variation in the signal types used in two genres that goes beyond graphical norms, it may be because those genres call for different relations to be used, rather than because the genre is signaling the same relations differently.

It is somewhat surprising that we see such limited variation in the signaling of individual relations

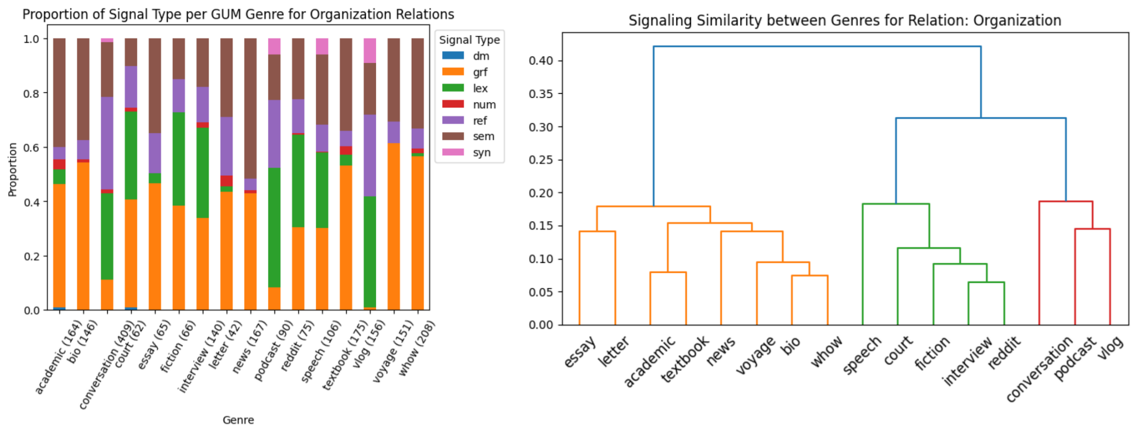

Figure 6: (Left) Proportions of relation signal types across genres for the `organization` relation. (Right) Dendrogram showing the signaling similarity between genres for the `organization` relation.

across genres, particularly considering that RST is a pragmatic formalism, and thus does not have restrictions on the structural components that must be present in order apply a specific discourse relation. Our results suggests that, despite being pragmatically defined, the discourse relations in the RST relation inventory display some degree of structural consistency in their manner of signaling. However, it is also worth noting that many of the signaling annotations from the GUM corpus which we are analyzing were automatically annotated by NLP tools/scripts. These automatic processes rely on restrictive heuristics, which may artificially limit the signaling variation being captured by the annotations. In future work, it would be beneficial to consider the specific limitations being imposed by such automatic annotations.

7 Conclusion

In this paper, we explored the cross-genre variation in how discourse relations are signaled in the GUM Corpus. We looked at the proportions of discourse signals in each genre, and we saw that there is a relative stability in how discourse relations are signaled across genres. We then conducted an analysis of which discourse relations display the most inter-genre variation in how they are signaled, using as a pairwise average of the JSD scores between different genres (Avg. Pairwise JSD) a metric of the inter-genre variability of individual discourse relations. We found that `organization`, `restatement`, and `explanation` relations display the most inter-genre variation, and that `evaluation` and `adversative` relations show the least inter-genre variation. Amongst the relations displaying the most inter-genre variation, we saw that the divide between spoken genres and written genres, and the accompanying divergence in graphical norms between the two modalities, is salient in accounting for the observed variation. Overall, we found that the RST discourse relations in GUM are signaled in a relatively stable manner across genres, and that the variation that does exist seem to largely come from differences in graphical norms, rather than differences in linguistic content.

Limitations

As noted in Section 4, using JDS as a metric for inter-genre variation relies on there being enough data to satisfy the assumption that the frequency of occurrence of signals is representative for the way that a relation is signaled in that genre. However, not all of the genres in the GUM corpus have the same number of documents, and for those with less documents, such as essay, which only has 5 documents, it is less sure that the assumption is sound. Still, the results from our correlation metrics in Section 4 in suggest that 5 documents is sufficient to give a reasonably stable ranking.

Additionally, as noted the Section 6, many of the signal annotations in the GUM corpus were automatically generated with NLP tools/scripts, which may limit the observable degree of inter-genre variation for relation signaling. A greater understanding of the inter-genre variation for relation signaling could be had from looking at a larger number of manual annotations, or by better accounting for the biases introduced by the automatic annotation tools.

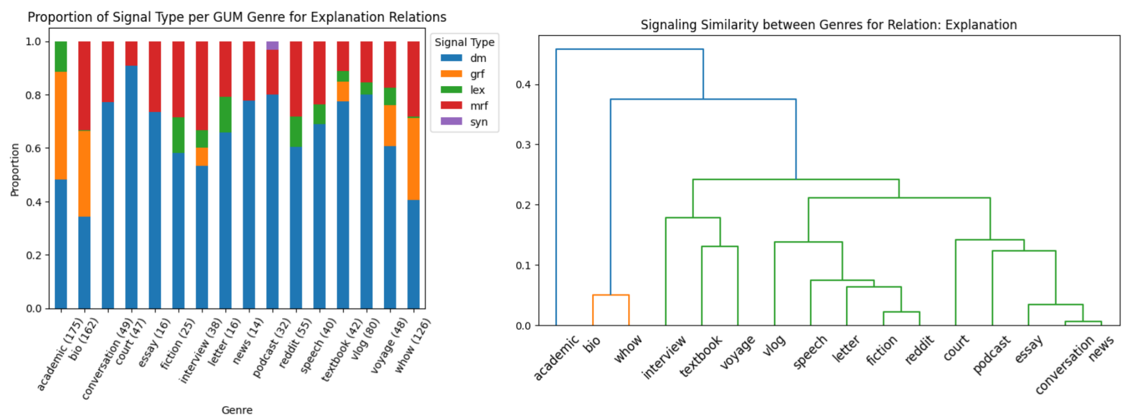

Figure 7: (Left) Proportions of relation signal types across genres for the `explanation` relation. (Right) Dendrogram showing the signaling similarity between genres for the `explanation` relation.

References

Lynn Carlson, Mary Ellen Okurowski, and Daniel Marcu. 2002. RST discourse treebank. *Linguistic Data Consortium, University of Pennsylvania.*

Debopam Das. 2019. Nuclearity in RST and signals of coherence relations. In *Proceedings of the Workshop on Discourse Relation Parsing and Treebanking 2019*, pages 30–37, Minneapolis, MN. Association for Computational Linguistics.

Debopam Das and Maite Taboada. 2018a. RST signalling corpus: A corpus of signals of coherence relations. *Language Resources and Evaluation*, 52:149–184.

Debopam Das and Maite Taboada. 2018b. Signalling of coherence relations in discourse, beyond discourse markers. *Discourse Processes*, 55(8):743–770.

Markus Egg and Debopam Das. 2022. Signalling conditional relations. *Linguistics Vanguard*, 8(s4):383–392.

Luke Gessler, Yang Janet 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.

Yang Liu and Amir Zeldes. 2019. Discourse relations and signaling information: Anchoring discourse signals in RST-DT. *Society for Computation in Linguistics*, 2(1).

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.

Maite Taboada and Debopam Das. 2013. Annotation upon annotation: Adding signalling information to a corpus of discourse relations. *Dialogue & Discourse*, 4(2):249–281.

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. 2024. eRST: A signaled graph theory of discourse relations and organization. *arXiv preprint arXiv:2403.13560*.

A RST Relation Inventory and eRST Signal Inventory in GUM v10

In Table 2 we include the RST relation inventory used in GUM v10, listing both coarse and fine-grained relations. For reference, in Figure 8 we include the eRST signaling inventory presented in Zeldes et al. (2024).

RST Relation Inventory			
Coarse	**Fine-grained**	**Coarse**	**Fine-grained**
ADVERSATIVE	ADVERSATIVE-ANTITHESIS ADVERSATIVE-CONCESSION ADVERSATIVE-CONTRAST	JOINT	JOINT-DISJUNCTION JOINT-LIST JOINT-SEQUENCE
ATTRIBUTION	ATTRIBUTION-POSITIVE		JOINT-OTHER
	ATTRIBUTION-NEGATIVE	MODE	MODE-MANNER
CAUSAL	CAUSAL-CAUSE		MODE-MEANS
	CAUSAL-RESULT	ORGANIZATION	ORGANIZATION-HEADING
CONTEXT	CONTEXT-BACKGROUND CONTEXT-CIRCUMSTANCE		ORGANIZATION-PHATIC ORGANIZATION-PREPARATION
CONTINGENCY	CONTINGENCY-CONDITION	PURPOSE	PURPOSE-ATTRIBUTE
ELABORATION	ELABORATION-ATTRIBUTE		PURPOSE-GOAL
	ELABORATION-ADDITIONAL	RESTATEMENT	RESTATEMENT-PARTIAL
EXPLANATION	EXPLANATION-EVIDENCE		RESTATEMENT-REPETITION
	EXPLANATION-JUSTIFY EXPLANATION-MOTIVATION	TOPIC	TOPIC-QUESTION TOPIC-SOLUTIONHOOD
EVALUATION	EVALUATION-COMMENT	SAME-UNIT	SAME-UNIT

Table 2: RST Relation Inventory in GUM v10.

signal type	subtypes	example
graphical	colon, dash, semicolon	*[Let me tell you a story* ***:****]*<organization–preparation>
	layout	*[Introduction]*<organization–heading>
	items in sequence	*1. wash [2. cut]*<joint–list>
	parentheses, quotation marks	*it rained* [***(****and snowed a bit****)***]<elaboration–additional>
	question mark	*[Did you****?****]*<topic–question> *No.*
lexical	alternate expression	*He agreed. [****That is*** *he said yes]*<restatement–repetition>
	indicative word/phrase	*They planned a party! [That's* ***nice****/****Can't wait!****]*<evaluation–comment>
morphological	mood	***Go*** *with them [I think you should]*<explanation–motivation>
	tense	*I* ***started*** *an hour ago, [now* ***I'm*** *resting]*<joint–sequence>
numerical	same count	*[****Two*** *reasons.]*<organization–preparation> *First. . .*
reference	comparative	*[I don't want* ***it****]*<adversative–antithesis> *I want* ***another one****.*
	demonstrative / personal	*They met* ***Kim****. [****This person*** */* ***she*** *was. . .]*<elaboration–additional>
	propositional	***They met Kim****. [****This encounter*** *was. . .]*<elaboration–additional>
semantic	antonymy	*Beer is* ***cheap****, [wine is* ***expensive****]*<adversative–contrast>
	attribution source	*[****Kim*** *said]*<attribution–positive> *they would*
	lexical chain	*it was* ***funny*** *[so they* ***laughed****]*<causal–result>
	meronymy	***The house*** *was big, [****the door*** *two meters tall]*<elaboration–additional>
	negation	*Kim danced, [Yun did****n't*** *dance]*<adversative–contrast>
	repetition/synonymy	*They met* ***Dr. Kim****. [****Dr. Kim****/****The surgeon*** *was. . .]*<elaboration–additional>
syntactic	infinitival/relative clause	*a plan [****to*** *win]*<purpose–attribute>
	interrupted matrix clause	*[I meant –]*<orgnization–phatic> *I mean,*
	modified head	*a* ***plan*** *[to win]*<purpose–attribute>
	nominal modifier	*articles [****explaining*** *chess]*<elaboration–attribute>
	parallel syntactic construction	***it's all*** *tasty [****it's all*** *pretty]*<joint–list>
	past/present participial clause	*Kim appeared [****dressed*** *in black]*<elaboration–attribute>
	reported speech	*[Kim said]*<attribution–positive> ***that*** *they* ***would***
	subject auxiliary inversion	*I would have [****had*** *I known]*<contingency–condition>

Figure 8: Signal inventory for eRST given in Zeldes et al. (2024): "Non-DM signal types and subtypes, with examples highlighting in red the signal tokens which indicate the relation of the unit in square brackets."