

From News to Medical: Cross-domain Discourse Segmentation

Elisa Ferracane¹, Titan Page², Junyi Jessy Li¹ and Katrin Erk¹

¹Department of Linguistics, The University of Texas at Austin

²Department of Linguistics, University of Colorado Boulder

elisa@ferracane.com, titan@colorado.edu

jessy@austin.utexas.edu, katrin.erk@mail.utexas.edu

Abstract

The first step in discourse analysis involves dividing a text into segments. We annotate the first high-quality small-scale medical corpus in English with discourse segments and analyze how well news-trained segmenters perform on this domain. While we expectedly find a drop in performance, the nature of the segmentation errors suggests some problems can be addressed earlier in the pipeline, while others would require expanding the corpus to a trainable size to learn the nuances of the medical domain.¹

1 Introduction

Dividing a text into units is the first step in analyzing a discourse. In the framework of Rhetorical Structure Theory (RST) (Mann and Thompson, 1988), the segments are termed elementary discourse units (EDUs), and a complete RST-style discourse analysis consists of building EDUs into a tree that spans the entire document. The tree edges are labeled with relations types, and nodes are categorized by their nuclearity (roughly, importance). RST segmentation is often regarded as a solved problem because automated segmenters achieve high performance (F1=94.3) on a task with high inter-annotator agreement (kappa=0.92) (Wang et al., 2018; Carlson et al., 2001). In fact, many RST parsers do not include a segmenter and simply evaluate on gold EDUs. However, numerous studies have shown errors in segmentation are a primary bottleneck for accurate discourse parsing (Soricut and Marcu, 2003; Fisher and Roark, 2007; Joty et al., 2015; Feng, 2015). Notably, even when using a top-performing segmenter, results degrade by 10% on the downstream tasks of span, nuclearity and relation labeling when using predicted instead of gold EDUs (Feng, 2015).

¹Code and data available at <http://github.com/elisaF/news-med-segmentation>.

[Patients were excluded] [if they had any other major Axis I psychiatric disorder, any medical or neurological disorder] [that could influence the diagnosis or treatment of depression,] [any condition other than depression] [that was not on stable treatment for at least the past one month,] [any condition] [that could pose a health risk during a clinical trial,] [and any clinically significant abnormality or disorder] [that was newly detected during the baseline assessments.]

Table 1: An example sentence from the novel MEDICAL corpus, with EDUs annotated in square brackets.

Separately, all available discourse segmenters are trained on news, and their ability to generalize to other domains, such as medical text, has not been well-studied. In our work, we focus on the medical domain because it has garnered cross-disciplinary research interest with wide-reaching applications. For example, the Biomedical Discourse Relation Bank was created for PDTB-style discourse parsing of biomedical texts (Prasad et al., 2011), and has been used to analyze author revisions and causal relations (Zhang et al., 2016; Marsi et al., 2014).

This work studies discourse segmentation in the medical domain. In particular, we: (1) seek to identify difficulties that news-trained segmenters have on medical; (2) investigate how features of the segmenter impact the type of errors seen in medical; and (3) examine the relationship between annotator agreement and segmenter performance for different types of medical data.

To this end, we present the first small-scale medical corpus in English, annotated by trained linguists (sample in Table 1). We evaluate this corpus with three RST segmenters, finding an expected gap in the medical domain. We perform a

detailed error analysis that shows medical-specific punctuation is the largest source of errors in the medical domain, followed by different word usage in syntactic constructions which are likely caused by news-derived word embeddings. Second, by comparing segmenters which use word embeddings versus syntax trees, we find access to parsed trees may not be helpful in reducing syntactically-resolvable errors, while an improved tokenizer would provide small benefits. Third, we note patterns between humans and segmenters where both perform better on extremely short texts and worse on those with more complex discourse.

We conclude with suggestions to improve the segmenter on the medical domain and recommendations for future annotation experiments.

Our contributions in this work are two-fold: a high-quality small-scale corpus of medical documents annotated with RST-style discourse segments; a quantitative and qualitative analysis of the discourse segmentation errors in the medical domain that lays the groundwork for understanding both the strengths and limits of existing RST segmenters, and the next concrete steps towards a better segmenter for the medical domain.

2 Related Work

Corpora in non-news domains. The seminal RST resource, the RST Discourse Treebank (RST-DT) (Carlson et al., 2001), consists of news articles in English. With the wide adoption of RST, corpora have expanded to other languages and domains. Several of these corpora include science-related texts, a domain that is closer to medical, but unfortunately also use segmentation guidelines that differ sometimes considerably from RST-DT² (research articles in Basque, Chinese, English, Russian, Spanish (Iruskieta et al., 2013; Cao et al., 2017; Zeldes, 2017; Yang and Li, 2018; Toldova et al., 2017; Da Cunha et al., 2012); encyclopedias and science news web pages in Dutch (Redeker et al., 2012)). Specifically in the medical domain, only two corpora exist, neither of which are in English. Da Cunha et al. (2012) annotate a small corpus of Spanish medical articles, and the RST Basque Treebank (Iruskieta et al., 2013) includes a small set of medical article abstracts. Our work aims to fill this gap by creating the first cor-

²A future direction of research could revisit this domain of science if the differing segmentation schemes are adequately resolved in the forthcoming shared task of the Discourse Relation Parsing and Treebanking 2019 workshop.

Corpus	#docs	#tokens	#sents	#EDUs
RST-DT SMALL	11	4009	159	403
MEDICAL	11	3356	169	399

Table 2: Corpus statistics.

pus of RST-segmented medical articles in English. Unlike several other works, we include all parts of the article, and not just the abstract.

Segmenters in non-news domains. While corpora have expanded to other domains, most automated discourse segmenters remain focused (and trained) on news. An exception is the segmenter in Braud et al. (2017a) which was trained on different domains for the purpose of developing a segmenter for under-resourced languages. However, they make the simplifying assumption that a single corpus represents a single (and distinct) domain, and do not include the medical domain. In this work, we study the viability of using news-trained segmenters on the medical domain.

3 Corpus Creation

Medical Corpus. The MEDICAL corpus consists of 2 clinical trial reports from PubMed Central, randomly selected for their shorter lengths for ease of annotation. We expect the language and discourse to be representative of this domain, despite the shorter length. As a result of the smaller size, we hypothesize annotator agreement and segmenter performance numbers may be somewhat inflated, but we nevertheless expect the nature of the errors to be the same. We divide the reports into their corresponding sections, treating each section as a separate document, resulting in 11 labeled documents. We chose to analyze sections individually instead of an entire report because moving to larger units typically yields arbitrary and uninformative analyses (Taboada and Mann, 2006). XML formatting was stripped, and figures and tables were removed. The sections for *Acknowledgements*, *Competing Interests*, and *Pre-publication History* were not included.

For comparison with the *News* domain, we created RST-DT-SMALL by sampling an equal number of Wall Street Journal articles from the “Test” portion of the RST-DT that were similar in length to the medical documents. The corpus statistics are summarized in Table 2.

Annotation Process. The annotation process was defined to establish a high-quality corpus that is

consistent with the gold-segmented RST-DT. Two annotators participated: a Linguistics graduate student (the first author), and a Linguistics undergraduate (the second author). To train on the task and to ensure consistency with RST-DT, the annotators first segmented portions of RST-DT. During this training phase, they also discussed annotation strategies and disagreements, and then consulted the gold labels. In the first phase of annotation on the medical data, the second author segmented all documents over a period of three months using the guidelines compiled for RST-DT (Carlson and Marcu, 2001) and with minimal guidance from the first author. In the second phase of annotation, all documents were re-segmented by both annotators, and disagreements were resolved by discussion.

Agreement. Annotators achieved on average a high level of agreement for identifying EDU boundaries with kappa=0.90 (averaged over 11 texts). However, we note that document length and complexity of the discourse influence this number. On a document of 35 tokens, the annotators exhibited perfect agreement. For the *Discussion* sections that make more use of discourse, the average agreement dropped to 0.84. The lowest agreement is 0.73 on a *Methods* section, which had more complex sentences with more coordinated sentences and clauses, relative clauses and nominal postmodifiers (as discussed in Section 6.1, these syntactic constructions are also a source of error for the automated segmenters).

4 Experiment

We automatically segment the documents in RST-DT SMALL and MEDICAL using three segmenters: (1) DPLP³ uses features from syntactic and dependency parses for a linear support vector classifier; (2) TWO-PASS (Feng and Hirst, 2014) is a CRF segmenter that derives features from syntax parses but also uses global features to perform a second pass of segmentation; (3) NEURAL (Wang et al., 2018) is a neural BiLSTM-CRF model that uses ELMo embeddings (Peters et al., 2018). We choose these segmenters because they are widely-used and publicly available (most RST parsers do not include a segmenter). DPLP has been cited in several works showing discourse helps on different NLP tasks (Bhatia et al., 2015). TWO-PASS, until recently, achieved SOTA on discourse segmentation when using parsed (not gold) syntax

³<https://github.com/jiyfeng/DPLP>

RST SEG	DOMAIN	F1	P	R
DPLP	<i>News</i>	82.56	81.75	83.37
	<i>Medical</i>	75.29	78.69	72.18
TWO-PASS	<i>News</i>	95.72	<u>97.19</u>	94.29
	<i>Medical</i>	84.69	86.23	83.21
NEURAL	<i>News</i>	<u>97.32</u>	95.68	<u>99.01</u>
	<i>Medical</i>	91.68	94.86	88.70

Table 3: F1, precision (P) and recall (R) of RST discourse segmenters on two domains (best numbers for *News* are underlined, for *Medical* are bolded).

trees. NEURAL now holds SOTA in RST discourse segmentation. We evaluate the segmenter’s ability to detect all EDU boundaries present in the gold data (not just intra-sentential) using the metrics of precision (P), recall (R) and F1.

The DPLP and TWO-PASS segmenters, both of which employ the Stanford Core NLP pipeline (Manning et al., 2014), were updated to use the same version of this software (2018-10-05).

5 Results

Table 3 lists our results on *News* and *Medical* for correctly identifying EDU boundaries using the three discourse segmenters. As expected, the *News* domain outperforms the *Medical* domain, regardless of which segmenter is used. In the case of the DPLP segmenter, the gap between the two domains is about 7.4 F1 points. Note that the performance of DPLP on *News* lags considerably behind the state of the art (-14.76 F1 points). When switching to the TWO-PASS segmenter, the performance on *News* increases dramatically (+13 F1 points). However, the performance on *Medical* increases by only 3.75 F1 points. Thus, large gains in *News* translate into only a small gain in *Medical*. The NEURAL segmenter achieves the best performance on *News* and is also able to more successfully close the gap on *Medical*, with only a 5.64 F1 difference, largely attributable to lower recall.

6 Error Analysis

We perform an error analysis to understand the segmentation differences between domains and between segmenters.

6.1 Error Types

We first group errors of the best-performing NEURAL segmenter into *error types*. Here we discuss

ERROR TYPE	PREDICTED	GOLD
amb. lexical cue	[our performance]since the buy - out makes it imperative]	[our performance since the buy - out makes it imperative]
infinitival “to”	[the auto giants will move quickly]to buy up stakes]	[the auto giants will move quickly to buy up stakes]
correct	[you attempt to seize assets]related to the crime]	[you attempt to seize assets related to the crime]
tokenization	[as identified in clinical trials.8-11]It]is noteworthy]	[as identified in clinical trials .][{ 8-11 }][It is noteworthy]
end emb. EDU	[Studies][confined to medical professionals have shown]	[Studies][confined to medical professionals][have shown]
punctuation	[the safety of placeboetine][(PB) hydrochloride capsules]	[the safety of placeboetine][(PB)][hydrochloride capsules]

Table 4: Examples of the most frequent segmentation error types with the erroneous EDU boundaries highlighted in red for *News* (top) and *Medical* (bottom) with predicted and gold EDU boundaries in square brackets (square brackets for citations are changed to curly brackets to avoid confusion). For *News*, the boundaries are inserted incorrectly (false positives) and for *Medical* they are omitted incorrectly (false negatives).

the most frequent types in each domain and give examples of each in Table 4 with the predicted and gold EDU boundaries.

ambiguous lexical cue Certain words (often discourse connectives) are strongly indicative of the beginning of an EDU, but are nonetheless ambiguous because of nuanced segmentation rules. In the Table 4 example, the discourse connective “since” typically signals the start of an EDU (e.g., in the RST discourse relations of *temporal* and *circumstance*), but is not a boundary in this case because there is no verbal element. Other problematic words include “that”, signalling relative clauses (often, but not always treated as embedded EDUs), and “and” which may indicate a coordinated sentence or clause (treated as a separate EDU) but also a coordinated verb phrase (not a separate EDU). Note this phenomenon is different from distinguishing between discourse vs. non-discourse usage of a word, or sense disambiguation of a discourse connective as studied in Pitler and Nenkova (2009).

infinitival “to” The syntactic construction of *to+verb* can act either as a verbal complement (treated as the same EDU) or a clausal complement (separate EDU). In the Table 4 example, the infinitival “to buy” is a complement of the verb “move” and should remain in the same EDU, but the segmenter incorrectly segmented it.

tokenization This error type covers cases where the tokenizer fails to detect token boundaries, specifically punctuation. These tokenization errors lead to downstream segmentation errors since punctuation marks, often markers of EDU boundaries, are entirely missed when mangled together with their neighboring tokens, as in ‘trials.[8-

11]It’ in Table 4.

punctuation This error occurs when parentheses and square brackets are successfully tokenized, but the segmenter fails to recognize them as EDU boundaries. This error is expected for square brackets, as they do not occur in RST-DT, but frequently appear in the *Medical* corpus for citations. It is not clear why the segmenter has difficulty with parentheses as in the Table 4 example “(PB)”, since they do occur in *News* and further almost invariably mark an EDU boundary.

end of embedded EDU An embedded EDU breaks up a larger EDU and is typically a relative clause or nominal postmodifier with a verbal element.⁴ While the segmenter is good at identifying the beginning of an embedded EDU, it often fails to detect the end. An embedded EDU such as the one listed in Table 4 can be clearly identified from a syntactic parse: the verbal element ‘have shown’ attaches to the subject ‘Studies’ and not the nominal postmodifier as predicted by the segmenter.

correct This category describes cases where we hypothesize the annotator made a mistake and the segmenter is correct. In the Table 4 example, the nominal postmodifier with non-finite clause “related to the crime” is an embedded EDU missed by annotators.

6.2 Errors between domains

In Figure 1, we compare the distribution of the most frequent error types in *News* (left) and the most frequent in *Medical* (right).

In *News* Figure 1a, the errors are mostly false positives where the segmenter incorrectly inserts

⁴For a more complete definition, see the tagging manual.




Figure 1: Most frequent segmentation error types by domain, using the best discourse segmenter.

boundaries before *ambiguous lexical cues*, and before *infinitival "to"* clauses (that are verbal complements). Interestingly, Braud et al. (2017a) found the tricky distinction of clausal vs. verbal infinitival complements to also be a large source of segmentation errors. These two error types also occur in *Medical*, though not as frequently, in part because the *to+verb* construction itself occurs less in the medical corpus. The third category of *correct* consists mostly of cases where the segmenter correctly identified an embedded EDU missed by annotators, illustrating both the difficulty of annotation even for experts and the usefulness of an automated segmenter for both in-domain and out-of-domain data since this error type is attested in both domains.

In *Medical* Figure 1b, we first note a stark contrast in distribution between the domains. The error types most frequent in *Medical* are hardly present in *News*; that is, errors in the *Medical* domain are often exclusive to this domain. The errors are mostly false negatives where the segmenter fails to detect boundaries around medical-specific use of *punctuation* marks, including square brackets for citations and parentheticals containing mathematical notations, which are entirely absent in *News*. The segmenter often misses the *end of embedded EDUs*, and more frequently than in *News*. The difference in this syntactically-identifiable error points to a gap in the embedding space for words signalling relative clauses and nominal postmodifiers. Given that ELMo embeddings have been shown to capture some syntax (Tenney et al., 2018), we recommend using PubMed-trained ELMo embeddings.⁵ One may further hypothesize that adding syntactic parses to the segmenter would help, which we explore in

⁵This option is viable once the MEDICAL corpus is expanded to a large enough size for training.




Figure 2: Distribution of the most frequent error types on *Medical* when using the NEURAL and TWO-PASS segmenter.

Section 6.3. The third error of *tokenization* occurs mainly around square brackets (citations), and this specific token never occurs in *News*.

6.3 Errors between segmenters

The rules of discourse segmentation rely heavily on syntax. Most discourse segmenters include syntax parse trees with the notable exception of the NEURAL segmenter. While this is the best-performing segmenter, we question whether it could be improved further if it had access to syntax trees. We probe this question by comparing the NEURAL segmentation errors with those found in TWO-PASS, which does use syntax trees.

Figure 2 illustrates the proportion of error types using the two segmenters. Although TWO-PASS makes use of syntax trees, the frequency of the syntactically-identifiable *end of embedded EDU* error type is only slightly lower. Because we do not have gold trees, it is also possible the news-trained parser performs poorly on medical and leads to downstream errors. We visually inspect the parse trees for these error cases and find the syntactic clause signaling the embedded EDU is correctly parsed in half the cases. Thus, bad parse trees contribute only partially to this error, and we suspect better trees may not provide much bene-

SECTION	KAPPA	F1	#TOKENS
Summary	1.00	100	35
Introduction	0.96	86.58	258
Results	0.93	91.74	354
Abstract	0.89	95.08	266
Methods	0.86	92.99	417
Discussion	0.84	89.03	365

Table 5: Average inter-annotator agreement per section, ordered from highest to lowest, the corresponding average F1 of the NEURAL segmenter, and number of tokens (there are 2 documents per section, except 1 for Summary).

fit. This finding is consistent with the little help dependency trees provided for cross-lingual discourse segmentation in Braud et al. (2017b).

We further note the tokenizer for TWO-PASS makes no errors on the medical data, but conversely has a higher proportion of *punctuation* errors. This pattern suggests improving the tokenizer of the NEURAL segmenter may simply shift errors from one type to another. To test this hypothesis, we use pre-tokenized text and find roughly half the errors do shift from one type to the other, but the other half is correctly labeled. That is, performance actually improves, but only slightly (F1=+0.36, P=+0.50, R=+0.24).

6.4 Errors between annotators and segmenters

Here we compare the level of annotator agreement with the performance of the NEURAL segmenter. In Table 5, we see that both humans and the model do well on extremely short texts (*Summary*). However, high agreement does not always translate to good performance. The *Introduction* section is straightforward for the annotators to segment, but this is also where most citations occur, causing the segmenter to perform more poorly. Earlier, we had noted the *Discussion* section was the hardest for annotators to label because of the more complex discourse. These more ambiguous syntactic constructions also pose a challenge for the segmenter, with lower performance than most other sections.

7 Next Steps

Based on our findings, we propose a set of next steps for RST discourse analysis in the medical domain. A much faster annotation process can be adopted by using the NEURAL segmenter as a first

pass. Annotators should skip extremely short documents and instead focus on the more challenging *Discussion* section. During training, we recommend using medical-specific word embeddings and preprocessing pipeline.⁶ Addressing even one of these issues may yield a multiplied effect on segmentation improvements as the *Medical* domain is by nature highly repetitive and formulaic.

However, a future avenue of research would be to first understand what impact these segmentation errors have on downstream tasks. For example, using RST trees generated by the lowest-performing DPLP parser nevertheless provides small gains to text categorization tasks such as sentiment analysis (Ji and Smith, 2017). On the other hand, understanding the verb form, which proved to be difficult in the *Medical* domain, has been shown to be useful in distinguishing text on experimental results from text describing more abstract concepts (such as background and introductory information), which may be a more relevant task than sentiment analysis (de Waard and Maat, 2012).

8 Conclusion

As a first step in understanding discourse differences between domains, we analyze the performance of three discourse segmenters on *News* and *Medical*. For this purpose, we create a first, small-scale corpus of segmented medical documents in English. All segmenters suffer a drop in performance on *Medical*, but this drop is smaller on the best *News* segmenter. An error analysis reveals difficulty in both domains for cases requiring a fine-grained syntactic analysis, as dictated by the RST-DT annotation guidelines. This finding suggests a need for either a clearer distinction in the guidelines, or more training examples for a model to learn to distinguish them. In the *Medical* domain, we find that differences in syntactic construction and formatting, including use of punctuation, account for most of the segmentation errors. We hypothesize these errors can be partly traced back to tokenizers and word embeddings also trained on *News*. We finally compare annotator agreement with segmenter performance and find both suffer in sections with more complex discourse. Based on our findings, we have proposed (Section 7) a set of next steps to expand the corpus and improve the segmenter.

⁶<https://allennlp.org/elmo>, <https://allenai.github.io/scispaacy>

Acknowledgments

We thank the anonymous reviewers for their helpful feedback. The first author was supported by the National Science Foundation Graduate Research Fellowship Program under Grant No. 2017247409. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

References

- Parminder Bhatia, Yangfeng Ji, and Jacob Eisenstein. 2015. [Better document-level sentiment analysis from rst discourse parsing](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 2212–2218. Association for Computational Linguistics.
- Chloé Braud, Ophélie Lacroix, and Anders Søgaard. 2017a. [Cross-lingual and cross-domain discourse segmentation of entire documents](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 237–243. Association for Computational Linguistics.
- Chloé Braud, Ophélie Lacroix, and Anders Søgaard. 2017b. [Does syntax help discourse segmentation? not so much](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2432–2442. Association for Computational Linguistics.
- Shuyuan Cao, Nianwen Xue, Iria da Cunha, Mikel Iruskieta, and Chuan Wang. 2017. [Discourse segmentation for building a rst chinese treebank](#). In *Proceedings of the 6th Workshop on Recent Advances in RST and Related Formalisms*, pages 73–81. Association for Computational Linguistics.
- Lynn Carlson and Daniel Marcu. 2001. Discourse tagging reference manual. *ISI Technical Report ISI-TR-545*, 54:56.
- Lynn Carlson, Daniel Marcu, and Mary Ellen Okurovsky. 2001. [Building a discourse-tagged corpus in the framework of rhetorical structure theory](#). In *Proceedings of the Second SIGdial Workshop on Discourse and Dialogue*, pages 1–10. Association for Computational Linguistics.
- Iria Da Cunha, Eric San Juan, Juan Manuel Torres-Moreno, Marina Lloberese, and Irene Castellóne. 2012. Diseg 1.0: The first system for spanish discourse segmentation. *Expert Systems with Applications*, 39(2):1671–1678.
- Vanessa Wei Feng and Graeme Hirst. 2014. [Two-pass Discourse Segmentation with Pairing and Global Features](#). *arXiv preprint arXiv:1407.8215*.
- Wei Vanessa Feng. 2015. *RST-style discourse parsing and its applications in discourse analysis*. Ph.D. thesis, University of Toronto (Canada).
- Seeger Fisher and Brian Roark. 2007. [The utility of parse-derived features for automatic discourse segmentation](#). In *Proceedings of the 45th Annual Meeting of the Association of Computational Linguistics*, pages 488–495. Association for Computational Linguistics.
- Mikel Iruskieta, Mara Jesus Aranzabe, A Diaz de Ilaraza, Itziar Gonzalez, Mikel Lersundi, and O Lopez de Lacalle. 2013. The RST Basque TreeBank: an online search interface to check rhetorical relations. In *Anais do IV Workshop A RST e os Estudos do Texto*, pages 40–49. Sociedade Brasileira de Computação.
- Yangfeng Ji and Noah A. Smith. 2017. [Neural Discourse Structure for Text Categorization](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 996–1005. Association for Computational Linguistics.
- Shafiq Joty, Giuseppe Carenini, and Raymond T. Ng. 2015. [CODRA: A Novel Discriminative Framework for Rhetorical Analysis](#). *Computational Linguistics*, 41(3):385–435.
- 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.
- Christopher D. Manning, Mihai Surdeanu, John Bauer, Jenny Finkel, Steven J. Bethard, and David McClosky. 2014. [The Stanford CoreNLP natural language processing toolkit](#). In *Association for Computational Linguistics (ACL) System Demonstrations*, pages 55–60.
- Erwin Marsi, Pinar Øztürk, Elias Aamot, Gleb Valerjevich Sizov, and Murat Van Ardelan. 2014. Towards text mining in climate science: Extraction of quantitative variables and their relations.
- Matthew Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. [Deep contextualized word representations](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, (Volume 1: Long Papers)*, pages 2227–2237. Association for Computational Linguistics.
- Emily Pitler and Ani Nenkova. 2009. [Using syntax to disambiguate explicit discourse connectives in text](#). In *Proceedings of Joint conference of the 47th Annual Meeting of the Association for Computational Linguistics and the 4th International Joint Conference on Natural Language Processing of the Asian Federation of Natural Language Processing*, pages 13–16. Association for Computational Linguistics.

- Rashmi Prasad, Susan McRoy, Nadya Frid, Aravind Joshi, and Hong Yu. 2011. The biomedical discourse relation bank. *BMC bioinformatics*, 12(1):188.
- Gisela Redeker, Ildikó Berzlánovich, Nynke van der Vliet, Gosse Bouma, and Markus Egg. 2012. [Multi-layer discourse annotation of a dutch text corpus](#). In *Proceedings of the Eighth International Conference on Language Resources and Evaluation*. European Language Resources Association.
- Radu Soricut and Daniel Marcu. 2003. [Sentence level discourse parsing using syntactic and lexical information](#). In *Proceedings of the 2003 Human Language Technology Conference of the North American Chapter of the Association for Computational Linguistics*, pages 228–235. Association for Computational Linguistics.
- Maite Taboada and William C Mann. 2006. Rhetorical structure theory: Looking back and moving ahead. *Discourse studies*, 8(3):423–459.
- Ian Tenney, Patrick Xia, Berlin Chen, Alex Wang, Adam Poliak, R Thomas McCoy, Najoung Kim, Benjamin Van Durme, Sam Bowman, Dipanjan Das, et al. 2018. What do you learn from context? probing for sentence structure in contextualized word representations.
- Svetlana Toldova, Dina Pisarevskaya, Margarita Ananyeva, Maria Kobozeva, Alexander Nasedkin, Sofia Nikiforova, Irina Pavlova, and Alexey Shelepov. 2017. [Rhetorical relations markers in russian rst treebank](#). In *Proceedings of the 6th Workshop on Recent Advances in RST and Related Formalisms*, pages 29–33. Association for Computational Linguistics.
- Anita de Waard and Henk Pander Maat. 2012. Verb form indicates discourse segment type in biological research papers: Experimental evidence. *Journal of English for Academic Purposes*, 11(4):357–366.
- Yizhong Wang, Sujian Li, and Jingfeng Yang. 2018. [Toward fast and accurate neural discourse segmentation](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 962–967. Association for Computational Linguistics.
- An Yang and Sujian Li. 2018. [Scidtb: Discourse dependency treebank for scientific abstracts](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 444–449. Association for Computational Linguistics.
- Amir Zeldes. 2017. The GUM corpus: creating multilayer resources in the classroom. *Language Resources and Evaluation*, 51(3):581–612.
- Fan Zhang, Diane Litman, and Katherine Forbes-Riley. 2016. [Inferring discourse relations from pdtb-style discourse labels for argumentative revision classification](#). In *Proceedings of the 26th International Conference on Computational Linguistics: Technical Papers*, pages 2615–2624. The COLING 2016 Organizing Committee.