

UC Davis

UC Davis Previously Published Works

Title

Fast Rhetorical Structure Theory Discourse Parsing

Permalink

<https://escholarship.org/uc/item/7kb2683v>

Authors

Heilman, Michael
Sagae, Kenji

Publication Date

2015-05-10

Peer reviewed

Fast Rhetorical Structure Theory Discourse Parsing

Michael Heilman

Educational Testing Service
Princeton, NJ, USA
mheilman@ets.org

Kenji Sagae

Institute for Creative Technologies
University of Southern California
Los Angeles, CA, USA
sagae@ict.usc.edu

1 Introduction

In recent years, There has been a variety of research on discourse parsing, particularly RST discourse parsing (Feng and Hirst, 2014; Li et al., 2014b; Ji and Eisenstein, 2014; Joty and Moschitti, 2014; Li et al., 2014a). Most of the recent work on RST parsing has focused on implementing new types of features or learning algorithms in order to improve accuracy, with relatively little focus on efficiency, robustness, or practical use. Also, most implementations are not widely available.

Here, we describe an RST segmentation and parsing system that adapts models and feature sets from various previous work, as described below. Its accuracy is near state-of-the-art, and it was developed to be fast, robust, and practical. For example, it can process short documents such as news articles or essays in less than a second.

The system is written in Python and is publicly available at <https://github.com/EducationalTestingService/discourse-parsing>.

2 Tasks and Data

We address two tasks in this work: discourse segmentation and discourse parsing. Discourse segmentation is the task of taking a sequence of word and punctuation tokens as input and identifying boundaries where new discourse units begin. Discourse parsing is the task of taking a sequence of discourse units and identifying relationships between them. In our case, the set of these relationships form a tree.

For both, we follow the conventions encoded in the RST Discourse Treebank (Carlson et al., 2002). Here, we give a brief overview of the corpus. See Carlson et al. (2001) for more information.

The treebank uses a representation where discourse is represented as a tree, with labels on nodes indicating relationships between siblings. Most RST relationships have a nucleus, expressing the core content, and a satellite that contributes additional information to the nucleus. Probably the simplest example is the “attribution” relationship: attributed (e.g., quoted) text is labeled as the nucleus, and text indicating the source of the attributed text is labeled as the satellite, with an “attribution” subcategorization.

The leaves of the RST trees are “elementary discourse units” (EDUs), which are contiguous spans of tokens roughly similar to independent clauses. Most branching in RST trees is binary, with one satellite and one nucleus, though there are some relations that have multiple nuclei and no satellite (e.g., lists).

The RST corpus consists of a training set of 347 documents and a test set of 38 documents. The texts in the RST treebank are a subset of those in the Penn Treebank (Marcus et al., 1993). For this reason, we retrained the syntactic parser used in our system, ZPar (Zhang and Clark, 2011), on the subset of the Penn Treebank WSJ sections 2 to 21 texts not present in the RST treebank.

For development of the system, we split the training set into a smaller subset for model estimation and a development validation set similar in size ($n = 40$) to the RST treebank test set.

3 Discourse Segmenter Description

In this section, we describe and evaluate the discourse segmentation component of the system. Our discourse segmenter is essentially a reimplementation of the baseline system from Xuan Bach et al. (2012). We do not implement their reranked model, which is more complex to implement and probably less efficient, and we use the ZPar parser (Zhang and Clark, 2011) for automatic syntactic parsing.

3.1 Segmenter Model and Features

Following Xuan Bach et al. (2012), we model RST as a tagging problem. Specifically, for each token in a sentence, the system predicts whether that token is the beginning of a new EDU or the continuation of an EDU. For this task, we use a conditional random field (Lafferty et al., 2001) model with ℓ_2 regularization, using the CRF++ implementation (<https://crfpp.googlecode.com>). Also, we assume that a new sentence always starts a new EDU, regardless of the CRF output.

The CRF uses simple word and POS features as well as syntactic features. The word and POS features are as follows (note that by “word”, we mean word or punctuation token):

- the lowercased form of the current word
- the part-of-speech (POS) of the current word

The syntactic features are based on automatic parses from ZPar (using a retrained model as discussed in §2). For each of the following nodes in the syntactic tree, there are two features, one for the nonterminal symbol and the head word (e.g., “VP, said”), and one for the nonterminal symbol and the head POS (e.g., “VP, VBD”). Note that these features will not be used for the last token in a sentence since there is no subsequent token.

- N_p : the first common ancestor of the current token and the subsequent word
- the subtree of N_p that contains the current word
- the subtree of N_p that contains the subsequent word
- the parent of N_p
- the right sibling of N_p

	P	R	F1
CRFSeg	91.0	87.2	89.0
Bach-etal-2012 (Base)	91.4	90.1	90.7
Bach-etal-2012 (Reranking)	91.5	90.4	91.0
our system	90.2	83.5	86.7
Human agreement	98.5	98.2	98.3

Table 1: Discourse segmentation performance in terms of percentages precision (“P”), recall (“R”), and F1 score (“F1”) for the “B-EDU” tag.

All of these features are extracted for the current word, the previous 2 words, and next 2 words in the sentence.

3.2 Segmenter Evaluation

Following Xuan Bach et al. (2012), we evaluate segmentation performance using the gold standard EDUs from the RST treebank test set, using the F1 score for the tag indicating the beginning of a new EDU (“B-EDU”). Since new sentences always begin new EDUs, we exclude the first tag in the output (always “B-EDU”) for each sentence. We first tuned the CRF regularization parameter using grid search on the split of the training set used for development evaluations, using a grid of powers of 2 ranging from 1/64 to 64.

The results are shown in Table 1. For comparison, we include previous results, including human-human agreement, reported by Xuan Bach et al. (2012), using syntax from the Stanford Parser (Klein and Manning, 2003a) (it is not clear from the paper what parsing model was used). The “CRFSeg” results are for the system from Hernault et al. (2010).

We are uncertain as to the cause for the observed differences in performance, though we hypothesize that the differences are at least partially due to differences in syntactic parsing, which is a key step in feature computation.

4 Discourse Parser Description

In this section, we describe our RST parser. It borrows extensively from previous work, especially Sagae (2009).¹

¹Note that we do not include Sagae (2009) in our evaluations since only within-sentence parsing performance was reported in that paper.

4.1 Shift-Reduce Approach

Following Sagae (2009) and Ji and Eisenstein (2014), we use an “arc standard” shift-reduce approach to RST discourse parsing.

4.2 Parsing Model

The parser maintains two primary data structures: a queue containing the EDUs in the document that have not been processed yet, and a stack of RST subtrees that will eventually be combined to form a complete tree.

Initially, the stack is empty and all EDUs are placed in the queue. Until a complete tree is found or no actions can be performed, the parser iteratively chooses to perform shift or reduce actions. The shift action creates a new subtree for the next EDU on the queue.

Reduce actions create new subtrees from the subtrees on the top of the stack. There are multiple types of reduce actions. First, there are unary or binary versions of reduce actions, depending on whether the top 1 or 2 items on the stack will be included as children in the subtree to be created. Second, there are versions for each of the nonterminal labels (e.g., “satellite:attribution”).

Following previous work, we collapse the full set of RST relations to 18 labels. Additionally, we binarize trees as described by Sagae and Lavie (2005).

Following Sagae (2009) and Ji and Eisenstein (2014), we treat the problem of selecting the best parsing action given the current parsing state (i.e., the stack and queue) as a classification problem. We use multi-class logistic regression with an ℓ_1 penalty, as implemented in the scikit-learn package, to estimate our classifier.

The parser supports beam search and k -best parsing, though we use simple greedy parsing (i.e., we set the beam size and k to 1) for the experiments described here.

4.3 Parsing Features

To select the next shift or reduce action, the parsing model considers a variety of lexical, syntactic, and positional features adapted from various previous work on RST discourse parsing, such as that of Sagae (2009) and the systems we compare to in §5. The features are as follows:

- the previous action (e.g., “binary reduce to satellite:attribution”)
- the nonterminal symbols of the n th subtree on the stack ($n = 0, 1, 2$), and their combinations
- the nonterminal symbols of the children of the n th subtree on the stack ($n = 0, 1, 2$)
- the lowercased words (and POS tags) for the tokens in the head EDU for the n th subtree on the stack ($n = 0, 1$) and the first EDU on the queue
- whether, for pairs of the top 3 stack subtrees and the 1st queue item, the distance (in EDU indices) between the EDU head is greater than n ($n = 1, 2, 3, 4$)
- whether, for pairs of the top 3 stack subtrees and the 1st queue item, the head EDUs are in the same sentence
- for the head EDUs of top 3 stack subtrees and the 1st queue item, the syntactic head word (lowercased), head POS, and the nonterminal symbol of the highest node in the subtree
- syntactic dominance features between pairs of the top 3 stack items and 1st queue item, similar to (Soricut and Marcu, 2003)
- for each of the first 3 stack items or 1st queue item, whether that item starts a new paragraph

5 Parsing Experiments

Following (Marcu, 2000, pp. 143–144) and other recent work, we evaluate our system according to the F1 score over labeled and unlabeled spans of discourse units in the RST treebank test set. This evaluation is analogous to the evalb bracket scoring program commonly used for constituency parsing (<http://nlp.cs.nyu.edu/evalb/>). For comparison with previous results, we use gold standard discourse segmentations (but automatic syntactic parses from ZPar).

We report F1 scores for agreement with the gold standard on unlabeled EDU spans (“span”), spans labeled only with nuclearity (“nuclearity”), and fully labeled spans that include relation information (“relation”).

We first tuned the ℓ_1 regularization parameter using grid search on the split of the training set used

	syntax	span	nuclearity	relation
our system	ZPar (retrained)	83.5	68.1	55.1
Li et al. (2014a)	Stanford	84.0	70.8	58.6
Joty et al. (2013)	Charniak (retrained)	82.5	68.4	55.7
Joty and Moschitti (2014)	Charniak (retrained)	–	–	57.3
Feng and Hirst (2014)	Stanford	85.7	71.0	58.2
Li et al. (2014b)	Penn Treebank	82.9	73.0	60.6
Ji and Eisenstein (2014)	MALT	81.6	71.0	61.8
Human agreement	–	88.7	77.7	65.8

Table 2: Test set discourse parsing performance in terms of F1 scores (%), using gold standard discourse segmentation. “syntax” indicates the source of POS tags and syntactic parse trees: “Stanford” refers to the Stanford parser (Klein and Manning, 2003b), “MALT” refers to Nivre and Marsi (2007), and “Charniak” refers to Charniak (2000).

	syntax	span	nuclearity	relation
our system	ZPar (retrained)	83.5	69.3	57.4
our system	PTB	84.7	71.2	59.4

Table 3: Development set discourse parsing performance in terms of F1 scores (%), using gold standard discourse segmentation. “syntax” indicates the source of POS tags and syntactic parse trees.

for development evaluations, using a grid of powers of 2 ranging from 1/16 to 16. We selected the setting that led to the highest F1 score for fully labeled spans (i.e., relation F1).

We compare to recently reported results from Ji and Eisenstein (2014) (their DPLP general +features model), Feng and Hirst (2014), Li et al. (2014b), Joty and Moschitti (2014), Li et al. (2014a), and Joty and Moschitti (2014).² The results are shown in Table 2. The human agreement statistics were originally reported by Ji and Eisenstein (2014). For each system, the table indicates the source of POS tags and syntactic parse trees (“Penn Treebank” means that gold standard Penn Treebank trees and tags were used).

We observe that our system is relatively close to the others in terms of F1 scores. We hypothesize that the differences in performance are at least partially due to differences in syntactic parsing.

²Joty and Moschitti (2014) and Joty and Moschitti (2014) do not explicitly state the source of syntactic parsers, but we infer from Joty et al. (2012) that the Charniak (2000) parser was used, with a model trained on a subset of the Penn Treebank that did not include the RST treebank test set.

5.1 The effect of automatic syntax parsing

In order to show the effect of using automatic parsing, we report performance on the development set (§2), using either gold standard syntax trees from the Penn Treebank or the automatic syntax trees from our retrained ZPar model (§2) for computing features. The F1 scores are shown in Table 3 (note that we are reporting results using the optimal settings from grid search on the development set).

It appears that the performance difference between using automatic rather than gold standard syntax is about 1 to 2 points of F1 score.

5.2 Parsing Speed

In this section, we evaluate the speed of the parser. Most previous papers on RST parsing do not report runtime experiments, and most systems are not widely available or easy to replicate.

Our parser uses a shift-reduce parsing algorithm that has a worst-case runtime that is linear in the number of EDUs. For comparison, Li et al. (2014b) employ a quadratic time maximum spanning tree parsing approach. The approach from Joty et al. (2013) also uses a polynomial runtime algorithm.

Other linear time parsers have been developed (Feng and Hirst, 2014; Ji and Eisenstein, 2014). However, feature computation can also be a per-

formance bottleneck. Feng and Hirst (2014) report an average parsing time of 10.71 seconds for RST treebank test set documents (and 5.52 seconds for a variant) on a system with “four duo-core 3.0 GHz processors”, not including time for preprocessing or discourse segmentation. In contrast, our system takes less than half a second per test set document on average (mean = 0.40, S.D. = 0.40, min. = 0.02, max. = 1.85 seconds) on a 2013 MacBook Pro with an i7-4850HQ CPU at 2.30 GHz. Of course, these performance measurements are not completely comparable since they were run on different hardware. The preprocessing (ZPar) and segmentation (§3.1) steps are also similarly fast.

6 Conclusion

In this paper, we have presented a fast shift-reduce RST discourse segmenter and parser. The parser achieves near state-of-the-art accuracy and processes Penn Treebank documents in less than a second, which is about an order of magnitude faster than recent results reported by Feng and Hirst (2014).

Acknowledgments

We would like to thank Dan Blanchard, Xinhao Wang, and Keelan Evanini for feedback about the paper. We would also like to thank Dan Blanchard, Diane Napolitano, Nitin Madnani, Aoife Cahill, Chong Min Lee, Michael Flor, and Keisuke Sakaguchi for initial help with and feedback about the implementation.

References

- Lynn Carlson, Daniel Marcu, and Mary Ellen Okurowski. 2001. Building a discourse-tagged corpus in the framework of rhetorical structure theory. In *Proceedings of the Second SIGdial Workshop on Discourse and Dialogue - Volume 16*, SIGDIAL '01, pages 1–10, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Lynn Carlson, Daniel Marcu, and Mary Ellen Okurowski. 2002. *RST Discourse Treebank LDC2002T07*. Philadelphia.
- Eugene Charniak. 2000. A maximum-entropy-inspired parser. 1st Meeting of the North American Chapter of the Association for Computational Linguistics.
- Vanessa Wei Feng and Graeme Hirst. 2014. A linear-time bottom-up discourse parser with constraints and post-editing. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 511–521, Baltimore, Maryland, June. Association for Computational Linguistics.
- Hugo Hernault, Danushka Bollegala, and Mitsuru Ishizuka. 2010. A sequential model for discourse segmentation. In *Proceedings of the 11th International Conference on Computational Linguistics and Intelligent Text Processing, CICLing'10*, pages 315–326, Berlin, Heidelberg. Springer-Verlag.
- Yangfeng Ji and Jacob Eisenstein. 2014. Representation learning for text-level discourse parsing. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 13–24, Baltimore, Maryland, June. Association for Computational Linguistics.
- Shafiq Joty and Alessandro Moschitti. 2014. Discriminative reranking of discourse parses using tree kernels. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2049–2060, Doha, Qatar, October. Association for Computational Linguistics.
- Shafiq Joty, Giuseppe Carenini, and Raymond Ng. 2012. A novel discriminative framework for sentence-level discourse analysis. In *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, pages 904–915, Jeju Island, Korea, July. Association for Computational Linguistics.
- Shafiq Joty, Giuseppe Carenini, Raymond Ng, and Yashar Mehdad. 2013. Combining intra- and multi-sentential rhetorical parsing for document-level discourse analysis. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 486–496, Sofia, Bulgaria, August. Association for Computational Linguistics.
- Dan Klein and Christopher D. Manning. 2003a. Accurate unlexicalized parsing. In *Proceedings of the 41st Annual Meeting of the Association for Computational Linguistics*, pages 423–430, Sapporo, Japan, July. Association for Computational Linguistics.
- Dan Klein and Christopher D Manning. 2003b. Fast exact inference with a factored model for natural language parsing. In S. Becker, S. Thrun, and K. Obermayer, editors, *Advances in Neural Information Processing Systems 15*, pages 3–10. MIT Press.
- John Lafferty, Andrew McCallum, and Fernando C. N. Pereira. 2001. Conditional random fields: Probabilistic models for segmenting and labeling sequence data. In *Proceedings of the 18th International Conference*

- on *Machine Learning 2001 (ICML 2001)*, pages 282–289.
- Jiwei Li, Rumeng Li, and Eduard Hovy. 2014a. Recursive deep models for discourse parsing. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2061–2069, Doha, Qatar, October. Association for Computational Linguistics.
- Sujian Li, Liang Wang, Ziqiang Cao, and Wenjie Li. 2014b. 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, June. Association for Computational Linguistics.
- Daniel Marcu. 2000. *The Theory and Practice of Discourse Parsing and Summarization*. MIT Press.
- Mitchell P. Marcus, Mary Ann Marcinkiewicz, and Beatrice Santorini. 1993. Building a large annotated corpus of english: The penn treebank. *Computational Linguistics*, 19(2):313–330.
- J. Hall J. Nilsson A. Chaney G. Eryigit S. Kbler S. Marinov Nivre, J. and E. Marsi. 2007. Maltparser: A language-independent system for data-driven dependency parsing. 13(2):95–135.
- Kenji Sagae and Alon Lavie. 2005. A classifier-based parser with linear run-time complexity. In *Proceedings of the Ninth International Workshop on Parsing Technology*, pages 125–132, Vancouver, British Columbia, October. Association for Computational Linguistics.
- Kenji Sagae. 2009. Analysis of discourse structure with syntactic dependencies and data-driven shift-reduce parsing. In *Proceedings of the 11th International Conference on Parsing Technologies (IWPT'09)*, pages 81–84, Paris, France, October. Association for Computational Linguistics.
- Radu Soricut and Daniel Marcu. 2003. Sentence level discourse parsing using syntactic and lexical information. In *Proceedings of the 2003 Conference of the North American Chapter of the Association for Computational Linguistics on Human Language Technology - Volume 1, NAACL '03*, pages 149–156, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Ngo Xuan Bach, Nguyen Le Minh, and Akira Shimazu. 2012. A reranking model for discourse segmentation using subtree features. In *Proceedings of the 13th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 160–168, Seoul, South Korea, July. Association for Computational Linguistics.
- Yue Zhang and Stephen Clark. 2011. Syntactic processing using the generalized perceptron and beam search. *Computational Linguistics*, 37(1).