

Multilingual Neural RST Discourse Parsing

Zhengyuan Liu, Ke Shi, Nancy F. Chen

Institute for Infocomm Research, A*STAR, Singapore
{liu.zhengyuan, shi.ke, nfychen}@i2r.a-star.edu.sg

Abstract

Text discourse parsing plays an important role in understanding information flow and argumentative structure in natural language. Previous research under the Rhetorical Structure Theory (RST) has mostly focused on inducing and evaluating models from the English treebank. However, the parsing tasks for other languages such as German, Dutch, and Portuguese are still challenging due to the shortage of annotated data. In this work, we investigate two approaches to establish a neural, cross-lingual discourse parser via: (1) utilizing multilingual vector representations; and (2) adopting segment-level translation of the source content. Experiment results show that both methods are effective even with limited training data, and achieve state-of-the-art performance on cross-lingual, document-level discourse parsing on all sub-tasks.

1 Introduction

Rhetorical Structure Theory (RST) (Mann and Thompson, 1988) is one of the most influential theories of discourse analysis, under which a document is represented by a hierarchical discourse tree. As shown in Figure 1a, the leaf nodes of an RST tree are text spans named Elementary Discourse Units (EDUs), and the EDUs are connected by rhetorical relations (e.g., *Cause*, *Contrast*) to form larger text spans until the entire document is included. The rhetorical relations are further categorized to *Nucleus* (core part) and *Satellite* (subordinate part) based on their relative importance. Thus, document-level discourse parsing consists of three sub-tasks: tree construction, nuclearity determination and relation classification. Moreover, downstream natural language processing tasks can benefit from RST-based structure-aware document analysis, such as summarization (Liu and Chen, 2019; Xu et al., 2020) and machine comprehension (Gao et al., 2020).

By utilizing various linguistic characteristics (e.g., N -gram and lexical features, syntactic and organizational features), statistical approaches have obtained substantial improvement on the English RST-DT benchmark (Sagae, 2009; Hernault et al., 2010; Joty et al., 2013; Li et al., 2014b; Heilman and Sagae, 2015). Recently, neural networks have been making inroads into discourse analysis frameworks, such as attention-based hierarchical encoding (Li et al., 2016) and integrating neural-based syntactic features into a transition-based parser (Yu et al., 2018). Lin et al. (2019) and their follow-up work (Liu et al., 2019) successfully explored encoder-decoder neural architectures on sentence-level discourse analysis, with a top-down parsing procedure.

Although discourse parsing has received much research attention and progress, the models are mainly optimized and evaluated in English. The main challenge is the shortage of annotated data, since manual annotation under the RST framework is labor-intensive and requires specialized linguistic knowledge. For instance, the most popular benchmark English RST-DT corpus (Carlson et al., 2002) only contains 385 samples, which is much smaller than those of other natural language processing tasks. The treebank size of other languages such as German (Stede and Neumann, 2014), Dutch (Redeker et al., 2012) and Basque (Iruskieta et al., 2013) are even more limited. Such limitations make it difficult to achieve

This work is licensed under a Creative Commons Attribution 4.0 International License. License details: <http://creativecommons.org/licenses/by/4.0/>.

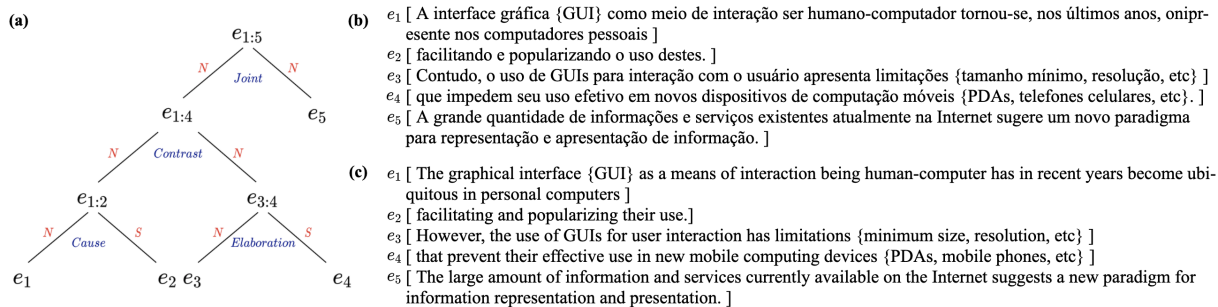


Figure 1: (a) An RST discourse tree example. e_i , $e_{j:k}$, N and S denote elementary discourse units, spans, nucleus and satellite respectively; (b) The original EDU segments in Portuguese; (c) The translated EDU segments in English.

acceptable performance on these languages required to fully support downstream tasks, and also lead to poor generalization ability of the computational approaches.

Since the treebanks of different languages share the same underlying linguistic theory, data-driven approaches can benefit from joint learning on multilingual RST resources (Braud et al., 2017). Therefore, in this paper, we investigate two methods to build a cross-lingual neural discourse parser: (1) From the embedding perspective: with the cross-lingual contextualized language models, we can train a parser on the shared semantic space from multilingual sources without employing a language indicator; (2) From the text perspective: since each EDU is a semantically-cohesive unit, we can unify the target language space by EDU-level translation, while preserving the original EDU segmentation and the discourse tree structures (see Figure 1c). To this end, we adapted and enhanced an end-to-end neural discourse parser, and investigated the two proposed approaches on 6 different languages. While the RST data for training is still in a small scale, we achieved the state-of-the-art performance on all fronts, significantly surpassing previous models, and even approaching the upper bound of human performance. Moreover, we conducted a topic modeling analysis on the collected multilingual treebanks to evaluate the model generality across various domains.

2 Methodology

2.1 Neural Discourse Parser

Since the encoder-decoder neural architecture with a top-down parsing procedure proposed in (Lin et al., 2019) has achieved impressive performance on sentence-level discourse analysis, here we adapted and enhanced it on the document-level parsing task. The neural model consists of an encoder, a span splitting decoder, and a nuclearity-relation classifier.

Encoder: The encoder produces EDU-level representations via a hierarchical encoding process. Given a document containing n tokens, an embedding layer is used to obtain token-level representations $\tilde{T} = \{\tilde{t}_1, \dots, \tilde{t}_n\}$. Then we obtain EDU-level representations by averaging the token embeddings for each EDU, and feed them to a Bi-GRU (Cho et al., 2014) component for document-level context-aware modeling. Moreover, to exploit implicit syntactic information like part-of-speech (POS) and sentential information (Yu et al., 2018), we incorporate boundary embeddings at both ends of each EDU from \tilde{T} to the context-aware vectors, and obtain the final EDU representation $E = \{e_1, \dots, e_m\}$, where m is the total EDU number.

Span Splitting Decoder: The decoder splits spans of EDUs to form the tree structure in a top-down transition-based procedure. Figure 2 illustrates the parsing steps of the example in Figure 1: the decoder maintains a *Stack*, which is initialized by the span of all EDUs $e_{1:m}$. At each decoding step t , the span $e_{i:j}$ at the head of stack is parsed into two sub-spans $e_{i:k}$ and $e_{k+1:j}$ ($i \leq k < j$), and k is the splitting position predicted via an attention-based pointer network (Bahdanau et al., 2015; Vinyals et al., 2015). Afterwards, spans containing more than one EDU are pushed into the stack, then the decoder iteratively parses the spans until *Stack* is empty.




Figure 2: Document-level neural parser. t , e and h denote input tokens, encoded EDU representations, and decoded hidden states. The *Stack* is maintained by the decoder to track top-down depth-first span splitting. With each splitting pointer k , sub-spans $e_{i:k}$ and $e_{k+1:j}$ are fed to a classifier Φ for nuclearity and relation determination.

Nuclearity-Relation Classifier: At each decoding step, after the span $e_{i:j}$ is split into two sub-spans $e_{i:k}$ and $e_{k+1:j}$, a bi-affine classifier (Dozat and Manning, 2017) is adopted to predict their nuclearity and relation labels. Here we use the joint labels of nuclearity and relation as previous studies (Yu et al., 2018; Lin et al., 2019). The total loss is specified as the sum of the cross entropy of span splitting and nuclearity-relation classification. Model implementation details and hyper-parameter configuration are described in Appendix A.

2.2 Multilingual Parsing

In this section, we introduce two approaches for the multilingual discourse parsing. Since both methods are model-independent, they can be adopted on various neural architectures, and extended to other low-resource scenarios.

2.2.1 Utilizing Cross-Lingual Vector Representations

Recently, the large-scale multilingual language models are able to provide universal encoders that project various inputs to a shared embedding space (Conneau and Lample, 2019), and are proved effective in natural language processing tasks such as machine comprehension. Therefore, to conduct discourse parsing on documents from various languages, we first propose to apply a cross-lingual representation backbone in the embedding layer in Section 2.1. Here, we utilize XLM-RoBERTa (Conneau et al., 2020), which supports 100 languages, and fine-tune it by joint training with the whole neural parser. Moreover, since BERT-based backbones usually have positional embedding limitation, to encode lengthy sequences without truncation for document-level discourse parsing, the sliding window strategy¹ is adopted for better long dependency modeling.

2.2.2 Adopting EDU Segment-Level Translation

Aside from using cross-lingual embedding, one alternative way is to transform multilingual text content into a monolingual space. While sophisticated neural approaches are able to generate multilingual translation with high quality and fluency, the commonly adopted sentence-level translation usually makes changes to the syntactic structure, which affects the original discourse annotation like the number and order of EDUs. Therefore, we propose to convert multilingual source content via EDU segment-level translation, as EDU segments are deemed to be semantically cohesive. We feed the documents with EDU segmentation (split by newlines) to a machine translator (Wu et al., 2016), then use the monolingual samples for training and evaluation. As shown in Figure 1, we observe that the translated material retains the original split and order at the EDU level, and shares the same English syntactic characteristics such as the position of discourse connective words (e.g., ‘however’, ‘although’) and relation pronouns (e.g., ‘that’, ‘which’). Then, we train the neural parser in Section 2.1 on the translated samples with their original tree structure, nuclearity, and relation annotations.

¹In our experiment, the input sequence length after sub-word tokenization can be larger than 2048. In order to exploit global contextual information, the window size is set as 500, and the stride size is 200.

Language	Tree Bank	Sample Num.	Relation Num.	Material Domain
English (En-DT)	English RST-DT (Carlson et al., 2002)	385	56	Wall Street Journal articles
	CST-News (Cardoso et al., 2011)	140	31	News articles
Brazilian	Summ-it (Collovini et al., 2007)	50	29	Science articles
Portuguese (Pt-DT)	Rhetalho (Pardo and Seno, 2005)	40	23	Articles in computer science
	CorpusTCC (Pardo and Nunes, 2004)	100	31	Articles in computer science and news
Spanish (Es-DT)	Spanish RST-DT (Da Cunha et al., 2011)	267	29	Text written by specialists on different topics (e.g. astrophysics, economy, law, linguistics)
German (De-DT)	MAZ Corpus (Stede and Neumann, 2014)	175	30	German newspaper commentaries
Dutch (Nl-DT)	Dutch RST-DT (Redeker et al., 2012)	80	31	Encyclopedias, letters, and news
Basque (Eu-DT)	Basque RST-DT (Iruskieta et al., 2013)	88	31	Abstracts from three specialized domains (medicine, terminology and science)

Table 1: The collected RST discourse treebanks from 6 languages.

Models	En-DT			Pt-DT			Es-DT			De-DT			Nl-DT			Eu-DT		
	Sp	Nu	Rel	Sp	Nu	Rel	Sp	Nu	Rel	Sp	Nu	Rel	Sp	Nu	Rel	Sp	Nu	Rel
MACRO F1 SCORE																		
Human*	88.7	77.7	65.8	-	78.0	66.0	86.0	82.5	76.8	-	-	-	83.0	77.0	70.0	81.7	-	61.5
MFS*	58.2	33.4	22.1	57.3	33.9	23.2	82.0	51.5	17.7	61.3	37.8	13.2	57.9	35.5	22.0	63.2	34.9	18.8
Li et al. (2014a)	85.0	70.8	58.6	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
Braud et al.(2017)	85.1	73.1	61.4	81.9	65.1	49.8	88.8	68.0	50.4	79.6	53.6	34.1	69.2	43.4	28.3	76.7	50.5	29.5
Cross-Lingual Representation																		
EN-Training	88.1	77.3	64.7	84.9	68.6	54.7	85.2	58.1	36.8	82.0	53.5	34.5	82.3	57.7	39.4	82.4	56.5	34.7
Multi-Training	88.9	77.5	65.7	87.2	72.9	60.8	88.3	75.3	60.5	84.1	62.8	45.9	86.4	64.6	49.5	85.9	67.3	51.9
Segment-Level Translation																		
EN-Training	88.3	77.8	64.9	85.1	69.2	55.0	85.9	58.6	37.0	80.7	53.3	34.1	82.6	58.0	39.7	82.6	58.3	36.2
Multi-Training	89.2	78.7	67.1	87.9	73.8	61.7	89.4	75.8	61.2	82.7	59.5	42.4	85.0	63.4	48.2	85.7	67.8	49.6
MICRO F1 SCORE																		
Yu et al. (2018)	85.5	73.1	60.2	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
Iruskieta (2019)	80.9	65.5	52.1	79.7	62.8	47.8	85.4	65.0	45.8	-	-	-	-	-	-	-	-	-
Cross-Lingual Representation																		
EN-Training	87.2	73.7	62.3	84.4	68.1	53.9	79.5	55.6	36.0	81.7	53.1	33.8	80.5	55.6	38.5	81.7	55.3	33.6
Multi-Training	87.5	74.7	63.0	86.3	71.7	60.0	86.2	71.1	54.4	83.6	62.2	45.1	85.9	64.5	49.4	85.1	65.8	47.7
Segment-Level Translation																		
EN-Training	87.4	74.6	62.8	84.9	68.0	54.2	82.6	56.3	35.1	79.5	52.7	34.0	81.5	57.0	39.1	81.2	57.3	35.5
Multi-Training	87.8	75.4	63.5	86.5	72.0	60.3	87.9	71.4	56.1	82.3	58.9	41.0	84.6	62.7	47.2	84.4	65.5	47.3

Table 2: Evaluation scores on multilingual RST treebanks. * denotes results from (Braud et al., 2017). *Sp*, *Nu* and *Rel* denote span splitting, nuclearity and relation determination respectively.

3 Experimental Results and Analyses

In this section, we describe data collection and present the experimental setting, results and analyses of the proposed methods.

3.1 Data and Pre-processing

We constructed a multilingual dataset by collecting treebanks from 6 languages: English (En-DT), Brazilian Portuguese (Pt-DT), Spanish (Es-DT), German (De-DT), Dutch (Nl-DT), and Basque (Eu-DT), and their details are shown in Table 1. Since the annotated formats are slightly different among treebanks, we conducted data pre-processing as in (Braud et al., 2017) to uniform them. All samples were transformed into binary trees, and units that were not linked to others within the tree were removed. Following (Lin et al., 2019), we reorganized the discourse relations to 18 categories, and attached the nuclearity labels *Nucleus-Satellite (NS)*, *Satellite-Nucleus (SN)* and *Nucleus-Nucleus (NN)* to the relation labels. For each language, we randomly selected 38 samples for evaluation. The total training set and test set are 1.1k and 228. For encoding input, we applied the pre-trained sub-word tokenizer of XLM-RoBERTa (Conneau et al., 2020). We adjusted random seeds to obtain multiple results for each language and used the average as reported scores.

3.2 Evaluation Result

The experimental results are shown in Table 2. Since macro-averaged and micro-averaged F1 scores are reported in different previous works, we conducted extensive comparisons using these two criteria.

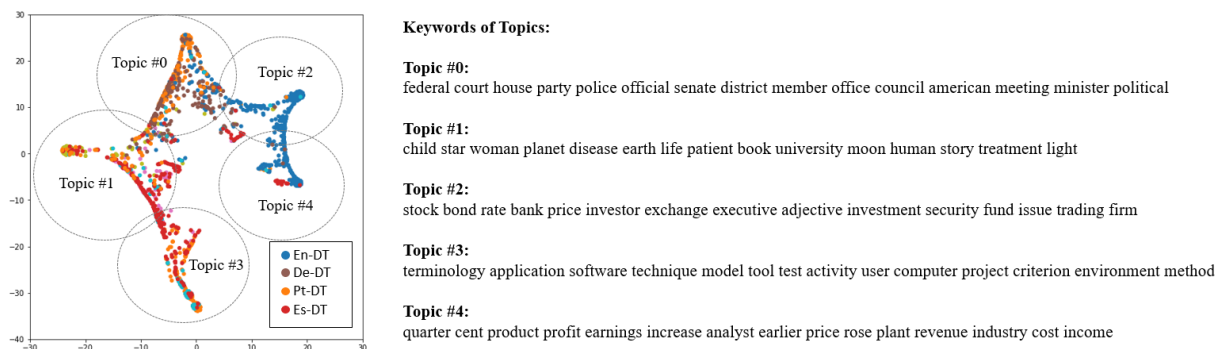


Figure 3: Analysis with topic modeling and feature visualization. Topic keywords are extracted with the LDA algorithm, and the scatter plot is illustrated by applying t-SNE.

The results demonstrate that (1) models which are only trained on the English treebank (*EN-Training*) can achieve competitive span splitting performance on the multilingual test sets; (2) the two proposed approaches with multilingual training (*Multi-Training*) surpass the baselines with a significant margin at all fronts: the span splitting prediction on all languages are approaching human performance, and nuclearity and relation determination are improved substantially compared to previously reported cross-lingual parsers (Braud et al., 2017); (3) Interestingly, the model with cross-lingual representation performs slightly better on the treebanks with fewer samples (e.g., De-DT, Ni-DT, and Eu-DT), and the model with segment-level translation obtains the best result in English.

3.3 Topic Modeling Analysis

To further assess the generality of our parsers from the domain perspective, we conducted a topic modeling analysis on the translated samples from multilingual treebanks. LDA (the topic number was set as 5) and t-SNE were used for topic modeling and feature visualization, respectively. As shown in Figure 3, the treebanks show a trend to cluster in different topics (marked in circles). For instance, the English treebank (En-DT) mainly focuses on the financial news domain (in blue). Compared to the Portuguese treebank (Pt-DT), the Spanish one (Es-DT) is more distinct to the En-DT, which is consistent with the performance gap between them under *EN-Training* (see Table 2). Therefore, by adding Spanish (Es-DT) and Portuguese (Pt-DT) data, topic coverage for the *Multi-Training* model expands to scientific and terminology articles, and thus becomes more generalizable to other domains.

4 Conclusion

In this paper, we investigated two approaches for cross-lingual neural discourse parsing. Experimental results show that both utilizing cross-lingual representation and adopting segment-level translation contribute to obtaining state-of-the-art performance on various treebanks. Moreover, monolingual models can also benefit from cross-lingual training by introducing data from more domains. For future work, we consider conducting domain adaption via few-shot learning to make our approach more generalizable.

Acknowledgements

This research was supported by funding from the Institute for Infocomm Research (I2R) under A*STAR ARES, Singapore. We thank N. Asher, A. T. Aw, S. Joty, W. Lei, B. Webber for insightful discussions, C. Braud for sharing linguistic resources, and the anonymous reviewers for their precious feedback to help improve and extend this piece of work.

References

Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural machine translation by jointly learning to align and translate. In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.

- Chloé Braud, Maximin Coavoux, and Anders Søgaard. 2017. Cross-lingual RST discourse parsing. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*, pages 292–304, Valencia, Spain, April. Association for Computational Linguistics.
- Paula CF Cardoso, Erick G Maziero, Mara Luca Castro Jorge, Eloize MR Seno, Ariani Di Felippo, Lucia Helena Machado Rino, Maria das Gracas Volpe Nunes, and Thiago AS Pardo. 2011. Cstnews-a discourse-annotated corpus for single and multi-document summarization of news texts in brazilian portuguese. In *Proceedings of the 3rd RST Brazilian Meeting*, pages 88–105.
- Lynn Carlson, Mary Ellen Okurowski, and Daniel Marcu. 2002. *RST discourse treebank*. Linguistic Data Consortium, University of Pennsylvania.
- Kyunghyun Cho, Bart van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. Learning phrase representations using RNN encoder–decoder for statistical machine translation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1724–1734, Doha, Qatar, October. Association for Computational Linguistics.
- Sandra Collovini, Thiago I Carbonel, Juliana Thiesen Fuchs, Jorge César Coelho, Lúcia Rino, and Renata Vieira. 2007. Summ-it: Um corpus anotado com informações discursivas visando a sumarização automática. *Proceedings of TIL*.
- Alexis Conneau and Guillaume Lample. 2019. Cross-lingual language model pretraining. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 7059–7069. Curran Associates, Inc.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. Unsupervised cross-lingual representation learning at scale. In *Proceedings of the Association for Computational Linguistics*. Association for Computational Linguistics.
- Iria Da Cunha, Juan-Manuel Torres-Moreno, and Gerardo Sierra. 2011. On the development of the rst spanish treebank. In *Proceedings of the 5th Linguistic Annotation Workshop*, pages 1–10.
- Timothy Dozat and Christopher D. Manning. 2017. Deep biaffine attention for neural dependency parsing. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*.
- Yifan Gao, Chien-Sheng Wu, Jingjing Li, Shafiq Joty, Steven CH Hoi, Caiming Xiong, Irwin King, and Michael R Lyu. 2020. Discern: Discourse-aware entailment reasoning network for conversational machine reading. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing*.
- Michael Heilman and Kenji Sagae. 2015. Fast rhetorical structure theory discourse parsing. *arXiv preprint arXiv:1505.02425*.
- Hugo Hernault, Helmut Prendinger, Mitsuru Ishizuka, et al. 2010. Hilda: A discourse parser using support vector machine classification. *Dialogue & Discourse*, 1(3).
- Mikel Iruskieta and Chloé Braud. 2019. EusDisParser: improving an under-resourced discourse parser with cross-lingual data. In *Proceedings of the Workshop on Discourse Relation Parsing and Treebanking 2019*, pages 62–71, Minneapolis, MN, June. Association for Computational Linguistics.
- Mikel Iruskieta, Maria J Aranzabe, Arantza Diaz de Ilarraza, Itziar Gonzalez, Mikel Lersundi, and Oier Lopez de Lacalle. 2013. The rst basque treebank: an online search interface to check rhetorical relations. In *4th workshop RST and discourse studies*, pages 40–49.
- 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.
- 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.

- Qi Li, Tianshi Li, and Baobao Chang. 2016. Discourse parsing with attention-based hierarchical neural networks. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 362–371.
- Xiang Lin, Shafiq Joty, Prathyusha Jwalapuram, and M Saiful Bari. 2019. A unified linear-time framework for sentence-level discourse parsing. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4190–4200, Florence, Italy, July. Association for Computational Linguistics.
- Zhengyuan Liu and Nancy Chen. 2019. Exploiting discourse-level segmentation for extractive summarization. In *Proceedings of the 2nd Workshop on New Frontiers in Summarization*, pages 116–121, Hong Kong, China, November. Association for Computational Linguistics.
- Linlin Liu, Xiang Lin, Shafiq Joty, Simeng Han, and Lidong Bing. 2019. Hierarchical pointer net parsing. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*, Hong Kong, China, November.
- 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.
- Thiago AS Pardo and Maria das Graças Volpe Nunes. 2004. Relações retóricas e seus marcadores superficiais: Análise de um corpus de textos científicos em português do brasil. *Relatório Técnico NILC*.
- Thiago Alexandre Salgueiro Pardo and Eloize Rossi Marques Seno. 2005. Rhetalho: um corpus de referência anotado retoricamente. *Anais do V Encontro de Corpora*, pages 24–25.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. 2019. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems*, pages 8024–8035.
- Gisela Redeker, Ildikó Berzlánovich, Nynke Van Der Vliet, Gosse Bouma, and Markus Egg. 2012. Multi-layer discourse annotation of a dutch text corpus. *age*, 1:2.
- 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*, pages 81–84. Association for Computational Linguistics.
- Manfred Stede and Arne Neumann. 2014. Potsdam commentary corpus 2.0: Annotation for discourse research. In *LREC*, pages 925–929.
- Oriol Vinyals, Meire Fortunato, and Navdeep Jaitly. 2015. Pointer networks. In *Advances in neural information processing systems*, pages 2692–2700.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, R’emi Louf, Morgan Funtowicz, and Jamie Brew. 2019. Huggingface’s transformers: State-of-the-art natural language processing. *ArXiv*, abs/1910.03771.
- Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, et al. 2016. Google’s neural machine translation system: Bridging the gap between human and machine translation.
- Jiacheng Xu, Zhe Gan, Yu Cheng, and Jingjing Liu. 2020. Discourse-aware neural extractive text summarization. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5021–5031, Online, July. Association for Computational Linguistics.
- Nan Yu, Meishan Zhang, and Guohong Fu. 2018. Transition-based neural rst parsing with implicit syntax features. In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 559–570.

A Model Description and Hyper-parameter Configuration

A.1 Details of Encoder-Decoder

Given a document containing n tokens $T = \{t_1, t_2, \dots, t_n\}$, the embedding layer (a pre-trained language model) produces token-level embeddings $\tilde{T} = \{\tilde{t}_1, \dots, \tilde{t}_n\}$, the EDU-level representations $C = \{c_1, \dots, c_m\}$ are calculated by averaging the respective token-level embeddings. Then, a multi-layer Bi-GRU is employed to generate the context-aware EDU-level representations $V = \{v_1, \dots, v_m\}$ by sequentially modeling the dependency among C , and each v_i is the concatenated vector of the the forward and the backward hidden states: $v_i = [\vec{v}_i; \overleftarrow{v}_i]$. Afterwards, the final EDU representations are produced via incorporating boundary embeddings at the beginning and end of each EDU from \tilde{T} to the context-aware EDU vector v_i :

$$e_i = W_e([v_i; \tilde{t}_{i.start}; \tilde{t}_{i.end}]) + b_e \quad (1)$$

where $;$ denotes the concatenation operation. W_e and b_e are the trainable parameter matrix and bias.

We employ a unidirectional GRU layer for the span splitting decoder, and its hidden state h_0 is initialized by the last hidden states of the encoder. At each decoding step, the hidden state h_t is produced by the GRU with the previous hidden state h_{t-1} and the input span representation $e_{i:k}$, where $e_{i:k}$ is calculated from taking the average of the respective EDU representations (*i.e.* $mean(e_i, \dots, e_k)$ for $e_{i:k}$). Then, the pointer network (Vinyals et al., 2015) is used to predict the splitting position according to the computed attention scores on encoded EDU representations, which is a softmax distribution over the input span.

$$s_{t,u} = \sigma(h_t, e_u) \quad \text{for } u = i \dots j \quad (2)$$

$$a_t = softmax(s_t) = \frac{exp(s_{t,u})}{\sum_{u=i}^j exp(s_{t,u})} \quad (3)$$

where $\sigma(x, y)$ is the dot product used as attention scoring function.

A.2 Details of Nuclearity-Relation Classifier

After decoder splits span $e_{i:j}$ into left sub-span $e_{i:k}$ and right sub-span $e_{k+1:j}$, the classifier first projects e_l and e_r to latent features \tilde{e}_l and \tilde{e}_r by a linear layer with Exponential Linear Unit (ELU), where e_l and e_r are the average of respective EDU representations in $e_{i:k}$ and $e_{k+1:j}$:

$$\tilde{e}_l = ELU(e_l^T U_1); \tilde{e}_r = ELU(e_r^T U_2) \quad (4)$$

Then a bi-affine layer with softmax activation maps the features to nuclearity-relation labels:

$$P_\theta(y|X) = softmax(\tilde{e}_l^T W_l + \tilde{e}_l^T W_{lr} \tilde{e}_r + \tilde{e}_r^T W_r + b) \quad (5)$$

where $W_l \in \mathcal{R}^{d \times R}$; $W_r \in \mathcal{R}^{d \times R}$ and $W_{lr} \in \mathcal{R}^{d \times d \times R}$ are the weights and bias $b \in \mathcal{R}^R$.

A.3 Training Loss

The parser's objective contains two parts: building the discourse tree structure and predicting the nuclearity and discourse relation labels. Therefore, the total loss is the sum of structure loss \mathcal{L}_s and label prediction loss \mathcal{L}_l , where \mathcal{L}_s is the cross entropy loss upon attention probabilities of the pointer network and \mathcal{L}_l is the cross entropy loss of the nuclearity-relation classification.

$$\mathcal{L}_s(\theta_s) = - \sum_{t=1}^T \log P_{\theta_s}(y_t | y_1, \dots, y_{t-1}, X) \quad (6)$$

$$\mathcal{L}_l(\theta_l) = - \sum_{m=1}^M \sum_{r=1}^R \log P_{\theta_l}(y_m = r | X) \quad (7)$$

where θ_s and θ_l are the parameters of the pointer network and classifier respectively, T is the total number of spans, and y_1, \dots, y_{t-1} denote the sub-trees that have been generated in the previous steps. M is the

number of spans that need to be split, and R is the number of nuclearity-relation labels. The total loss with L_2 -regularization is:

$$\mathcal{L}_{total}(\theta^*) = \mathcal{L}_s(\theta_s) + \mathcal{L}_l(\theta_l) + \lambda \|\theta^*\|_2^2 \quad (8)$$

where λ is the regularization strength and θ^* refers to all learning parameters of the model.

A.4 Hyper-parameter Configuration

The neural model was implemented in PyTorch (Paszke et al., 2019). We used ‘xlm-roberta-base’ implemented in (Wolf et al., 2019) and fine-tuned the last 4 layers during training. In order to exploit global contextual information, the window size was set as 500 and the stride size was 200. Documents were tokenized via the sub-word scheme as in (Conneau and Lample, 2019). We trained the model for 30 epochs and selected the best checkpoints on a validation set for evaluation. Adam optimization algorithm was used with batch size of 3, weight decay of $5e-5$, and learning rate of $1e-4$. Dropout rate was set as 0.5 during training. The embedding dimension and hidden size were 768 and 384. The trainable parameter size was 67M, where 31M parameters were from fine-tuning the language model. All experiments were run on a Tesla V100 GPU with 16GB memory.