

June 14–16, 2023

Light Coreference Resolution for Russian with Hierarchical Discourse Features

Elena Chistova and Ivan Smirnov

FRC CSC RAS

Moscow, Russia

{chistova, ivs}@isa.ru

Abstract

Coreference resolution is the task of identifying and grouping mentions referring to the same real-world entity. Previous neural models have mainly focused on learning span representations and pairwise scores for coreference decisions. However, current methods do not explicitly capture the referential choice in the hierarchical discourse, an important factor in coreference resolution. In this study, we propose a new approach that incorporates rhetorical information into neural coreference resolution models. We collect rhetorical features from automated discourse parses and examine their impact. As a base model, we implement an end-to-end span-based coreference resolver using a partially fine-tuned multilingual entity-aware language model LUKE. We evaluate our method on the RuCoCo-23 Shared Task for coreference resolution in Russian. Our best model employing rhetorical distance between mentions has ranked 1st on the development set (74.6% F1) and 2nd on the test set (73.3% F1) of the Shared Task¹. We hope that our work will inspire further research on incorporating discourse information in neural coreference resolution models.

Keywords: coreference resolution, Rhetorical Structure Theory, referential choice, rhetorical distance, Russian

DOI: 10.28995/2075-7182-2023-22-34-41

Разрешение кореференции для русского языка с использованием признаков иерархического дискурса

Чистова Е. В., Смирнов И. В.

ФИЦ ИУ РАН

Москва, Россия

{chistova, ivs}@isa.ru

Аннотация

Разрешение кореференции – это задача выявления и группировки упоминаний, относящихся к одному и тому же объекту реального мира. При решении задачи методами глубокого обучения в первую очередь обращают внимание на проблемы обучения векторных представлений сущностей и оценки вероятности наличия кореферентной связи между ними. Однако существующие методы не позволяют в явном виде учитывать референциальный выбор в иерархическом дискурсе. В данной работе оценивается важность признаков, полученных на основе автоматического риторического анализа, применительно к нейросетевым моделям. В качестве базового метода реализована end-to-end архитектура с использованием мультязычной языковой модели LUKE, учитывающей при кодировании текста границы сущностей. Лучшая модель, в которой используется признак риторического расстояния между сущностями, занимает первое место на валидационной (74.6% F1) и второе место на тестовой (73.3% F1) выборке соревнования RuCoCo-2023.

Ключевые слова: разрешение кореференции, теория риторических структур, референциальный выбор, риторическое расстояние, русский язык

¹The code and models are available at <https://github.com/tchewik/corefhd>

1 Introduction

Coreference resolution is the task of identifying and grouping mentions referring to the same real-world entity. It is a challenging task in natural language processing, as it often requires both linguistic and common knowledge. In recent years, neural models have achieved remarkable success in coreference resolution. These models aim to identify mention spans and assign pairwise scores. However, they mostly rely on surface explicit features, such as the distance between entities in tokens, and overlook the hierarchical discourse structure. Contextual word embeddings, despite their morphosyntactic and semantic richness, also have limitations in capturing document discourse beyond local cues.

Our system for RuCoCo-2023, called CorefHD (**C**oreference in **H**ierarchical **D**iscourse), enhances the classical neural architecture with automatically retrieved features that capture aspects of hierarchical discourse. It uses pretrained transformer-based contextualized word embeddings, along with dense embeddings of hierarchical discourse features: linear distance, rhetorical distance, and anaphor-to-LCA distance. To retrieve the discourse hierarchy of the text, we use an RST parser predicting constituency trees in accordance with the Rhetorical Structure Theory (Mann and Thompson, 1988).

The main contributions of this paper are:

- We propose a new method that incorporates discourse information into neural coreference resolution models.
- We test various discourse features that capture the distances between mentions on a large coreference resolution dataset in Russian.
- We apply a number of memory reduction techniques and demonstrate that high-quality coreference resolution can be done with standard neural architecture even with limited computational resources.
- We use a multilingual entity-aware LUKE (Yamada et al., 2020) language model and show that it performs competitively with the monolingual language models for Russian in coreference resolution.
- We join the RuCoCo-2023 Shared Task, and achieve 1st place on the development set and 2nd place on the test set of the contest with the model using the rhetorical distance feature.

The rest of this work is organized as follows: Section 2 reviews a concept of referential distance and current work on coreference resolution in hierarchical discourse. Section 3 describes our method in detail. Section 4 presents our experimental setup. Section 5 analyzes our results. Section 6 concludes the paper and discusses future work.

2 Related Work

Linear referential distance measures how many clauses separate an anaphor from its antecedent (Givón, 1983). However, not all phrases in discourse require the same level of attention. It is observed (Grosz and Sidner, 1986) that the discourse structure of a text contains discourse units inside and outside the intention and attention. Using a corpus of 30 manually annotated texts, it is shown (Cristea et al., 1999) that a hierarchical model of discourse has greater potential for improving the coreference resolution performance than a linear model of discourse. The most popular hierarchical discourse framework as of today is Rhetorical Structure Theory (Mann and Thompson, 1988). Within RST, one can consider in the referential distance the rhetorical structures, where attention focus is part of the definition (Moser and Moore, 1996) of subordinating (mononuclear) RST relations. An approach to computing referential distance with respect to the rhetorical tree is suggested by Kibrik (Kibrik, 1999): the rhetorical distance can be measured by counting the nodes in an RST tree that are visited while walking from the mention to its possible antecedent. A study on the RST Discourse Treebank² shows that while rhetorical distance does not imply the one and only referential choice, it is still one of the principal factors for referential choice prediction (Kibrik and Krasavina, 2005). Another study (Fedorova et al., 2010) uses six RST-annotated text fragments in Russian to demonstrate that rhetorical distance has a significant impact on the referent activation in working memory.

Closest to our work are (Khosla et al., 2021) implementing various features over an RST tree produced with a parser for English. However, their main concern is how general is the lowest common ancestor of

²<https://catalog.ldc.upenn.edu/LDC2002T07>

two mentions in the rhetorical constituency tree. While this is somewhat related to the working memory load of keeping two mentions active, they do not directly consider a concept of referential distance and, most importantly, ignore nuclearity (i.e. attention), which is a crucial feature in rhetorical structures.

In this paper, we apply the RST parser for Russian to build hierarchical discourse trees. The distance features obtained from these trees we use in a neural coreference resolution model. As far as we know, we are the first to model referential distances in hierarchical discourse with neural models. We also examine the impact of the RST features in coreference resolution for Russian on a large annotated corpus.

3 Approach

End-to-end coreference resolution involves finding entities in plain text and collecting them into clusters so that each cluster corresponds to a single real-world object. As a core method, we apply the classical (Lee et al., 2018)’s approach to end-to-end coreference parsing with a span-ranking architecture, except for the higher-order inference which has been proven to be ineffective (Xu and Choi, 2020). This approach to coreference resolution involves five main steps:

1. Collect the initial set of spans.
2. Rank the collected spans with a linear transformation of span embeddings and keep the top-k resembling entities.
3. Collect the coarse referent-to-antecedent probabilities for each possible pair of entities. This is calculated as a sum of corresponding span probabilities obtained in the previous step and a score obtained with a bilinear transformation of two *mention encodings*. Keep the top-n pairs with the highest prediction.
4. Compute the final coreference scores for each possible mention-antecedent pair that made it to this step. This is done with a feedforward layer processing *mention pair encodings*. Assign to each mention the antecedent with the highest predicted probability.
5. The predictions form connected chains of mentions that can be viewed as clusters.

The following gives the details of how our system encodes entities and their pairs.

Mention Encoding Each fine-grained token is encoded as an average of its subtoken representations obtained using a language model. The initial entity candidates are collected greedily, with the only parameter being the maximum length of the span. To adjust this parameter effectively, we use token representations instead of LM subtoken representations. Since language models work with a limited context, we collect each paragraph representation separately.

Mention Pair Encoding To calculate the final predictions for each pair of found mentions, we use a feedforward layer that takes a mention pair embedding as input. This embedding consists of the concatenation of two individual mention encodings and the embedding of the token count between them. For the models employing discourse hierarchy features, we represent them similarly to token distances and concatenate them to the pair embeddings.

3.1 Discourse Hierarchy Features

Given two spans i (a mention) and j (its possible antecedent), we first find the elementary discourse units u_i and u_j covering the corresponding spans in a predicted RST tree. Then we compute the discourse-related features and concatenate them with mention pair encoding.

Two metrics are used to measure referential distance in discourse, as outlined in (Kibrik, 1999):

- **Linear Distance** (D_{Lin}) in our model is a number of predicted elementary discourse units (EDUs) occurring between two spans.
- **Rhetorical Distance** (D_{Rh}) is a number of nuclear EDUs occurring between two spans in a hierarchical rhetorical tree.

We also adopt a feature estimating the amount of generality required to have two mentions in the same discourse subtree (Khosla et al., 2021):

- **Referent’s distance to the LCA** (D_{LCA}) Assuming mention i always appearing to the right of any possible antecedent j , and $LCA(u_j, u_i)$ being the lowest discourse unit covering both u_i and u_j in the constituency RST tree, $D_{LCA} = \text{dist}(u_i, LCA(u_j, u_i))$.

4 Experimental Setup

4.1 Pretrained Language Model

We employ the multilingual LUKE³ (Ri et al., 2022). It is a language model that has been trained with both masked language modeling (MLM) and masked entity prediction (MEP) tasks. The entity annotations in the training corpus are collected from hyperlinks in Wikipedia dumps. This multilingual model has previously demonstrated significant improvement in question answering and cloze prompt tasks for Russian compared to mBERT (Devlin et al., 2019) and XLM-RoBERTa (Conneau and Lample, 2019). We hypothesize that explicit coreference resolution can also benefit from LM-ingrained entity encoding.

4.2 Factors Reducing Memory Consumption

Neural coreference resolution is a memory-intensive task. The common approach to end-to-end coreference resolution (Lee et al., 2017; Lee et al., 2018) requires computation over each and every span in a document. A number of recent works suggest more optimal alternative methods, in which the object of processing is not a span but a token (Kirstain et al., 2021; Thirukovalluru et al., 2021; Dobrovolskii, 2021). Despite this, the relevant research adopting language model fine-tuning still requires 40 to 80 GB of video memory (Dobrovolskii, 2021; Mæhlum et al., 2022). In our study, we investigate the extent to which the most classical span-based approach to coreference resolution can be scaled down.

Each our model is trained on a single 32GiB Tesla V100 GPU, with peak memory allocation of 98%. To achieve this, we modified the standard model architecture and implementation:

- The main factor that allows a coreference model to be trained on a large dataset with limited memory is *excluding full LM fine-tuning*. In our experiments, a language model is frozen except for the last k layers. The value of k is determined empirically by the amount of video memory available. In our setting, $k = 8$ of 23 layers.
- After initial token encoding, the bidirectional LSTM is used to obtain *lower-dimensional token embeddings*. The span embedding is computed from the sequence of compressed token embeddings using self-attention. In our experiments, $\mathbf{e}_{LM} \in \mathbb{R}^{1024}$ and $\mathbf{e}_{LSTM} \in \mathbb{R}^{100}$.
- *Each paragraph of the text is encoded with a language model separately*. This allows long news articles to be encoded without trimming and high-dimensional partially-trainable LM embeddings to be compressed in place, thereby saving memory.

We also use standard techniques reducing memory requirements:

- *Batch size = 1*. Gradient accumulation did not improve training results.
- All the calculations are performed with *mixed precision*.

4.3 Instruments for Linguistic Analysis

Tokenization and sentence splitting are performed with the Razdel⁴ library. Named entities are recognized with the SpaCy⁵ ru_core_news_lg model predicting BIO-tags from token embeddings. Discourse structures are produced with the IsaNLP RST⁶ parser for Russian (Chistova et al., 2021). The parser generates trees for each paragraph; we merged these trees with a right-branching multinuclear JOINT relation to construct the full-text RST trees.

4.4 Data

We perform the experiments on the RuCoCo-2023 Shared Task dataset described in (Dobrovolskii et al., 2022). It is a large corpus for coreference resolution collected from news articles in Russian. It contains annotated news in multiple categories, including finance, world news, sports, and more. The corpus

³studio-ousia/mluke-large-lite

⁴<https://github.com/natasha/razdel>

⁵<https://spacy.io/>

⁶https://github.com/tchewik/isanlp_rst

Original	Translation
Обитатели небоскребов Нью-Йорка спешат обзавестись [парашютами] _{SA}	Residents of New York skyscrapers rush to get [parachutes] _{SA}
Обитатели небоскребов Нью-Йорка спешат обзавестись [парашютами] _{SA} . Это связано с недавними терактами в этом городе.	Residents of skyscrapers in New York rush to get [parachutes] _{SA} . This is due to recent terrorist attacks in the city.
Одна из американских фирм по [их] _{SA} производству сообщила, что в офисе не прекращают звонить телефоны. Владельцы квартир в высотных зданиях интересуются возможностью приобретения [новой модели парашюта] _I , [которая] _I была разработана после трагических событий 11 сентября. [Он] _I стоит около 800 долларов и раскрывается автоматически. [...]	One of [their] _{SA} manufacturer reports that the phones in its office never stop ringing. Apartment owners in high-rise buildings are interested in buying [a new parachute] _I [which] _I is developed after the tragic events of September 11. [It] _I costs about \$800 and opens automatically. [...]

Table 1: Split-antecedent annotation example in the RuCoCo dataset, from 2001_world_new_003.

includes both single one-to-one coreference annotation and split antecedents one-to-many coreference annotation. However, the distinguishing feature of the latter is that it is annotated among clusters (entities), not mentions (an example is shown in Table 1). It poses a challenge in identifying pairs of mentions from different groups that are connected by split-antecedent relations. To address this additional challenge, our model’s architecture would require additional modifications. Although both tasks are evaluated jointly in the competition, this study’s emphasis is on the standard coreference resolution. Here, we conduct some additional analyses of the data relevant to our methods.

Firstly, it is critical for our model to determine the maximum entity length in the corpus. The results on the train set are illustrated in Fig. 1. The mean entity length is 2, and the maximum is 42. The maximum mention length in our system is set to 13 tokens, which covers 99.7% of entities in the corpus.




Figure 1: Entity lengths in tokens (train set).




Figure 2: Number of paragraphs per annotation.

Secondly, we examine the number of paragraphs in the data. It will be identical to the number of trees in RST parser output. Thus, if we construct the text-level tree by merging paragraph trees, it could be critical for long discourse dependencies. The results are shown in Fig. 2. Every line split is considered a paragraph. The median paragraph count is 9, with the maximum number of separated lines being 162. Some news articles are exceptionally long, and some of them include enumerated lists. Combining multiple trees into one can affect the referential distance estimation in a few particularly long texts.

4.5 Evaluation

In the Shared Task, the coreference resolution F1 score is calculated using the Link-based Entity Aware (LEA) metric (Moosavi and Strube, 2016). In this metric, the weight of each entity is determined by its size, with larger entities being considered more important. It also evaluates resolved coreference relations instead of resolved mentions.

The models are validated during training using 5% of the official train set. We run random splitting 4 times and report the average result. The listed results on the official development and test sets of the competition are obtained with the exact same models.

5 Results and Discussion

In Table 2, we present the results of our experiments on the development set of the RuCoCo-2023 Shared Task. We also report the performance of our system on the test set (also called the final set) of the RuCoCo-2023 Shared Task in Table 3. Our baseline model noticeably outperforms the RuRoBERTa-large-based baseline provided by the organizers, which achieved 68.4% and 67.4% F1 on the development and test sets, respectively.

	Precision	Recall	F1	Top-1 F1 (leaderboard)
Baseline	78.7 ± 0.7	69.1 ± 0.7	73.5 ± 0.5	74.3
+ D_{Lin}	78.6 ± 1.8	68.3 ± 2.2	73.0 ± 0.5	74.0
+ D_{Rh}	78.5 ± 1.5	69.3 ± 1.0	73.6 ± 0.9	74.6
+ D_{LCA}	75.0 ± 0.8	70.9 ± 1.0	72.9 ± 0.4	73.5

Table 2: Models evaluation on the official development set.

Due to the strict limit on the number of submissions in the final phase of the competition, we could only evaluate the two best performing models, Baseline and Baseline+ D_{Rh} , on a private leaderboard.

	Precision	Recall	F1	Top-1 F1 (leaderboard)
Baseline	79.1 ± 0.8	66.9 ± 0.6	72.5 ± 0.3	72.8
+ D_{Rh}	79.3 ± 1.6	66.6 ± 1.9	72.4 ± 0.5	73.3

Table 3: Models evaluation on the official test set (“Final”).

Features D_{Lin} and D_{LCA} are not found to be effective for the task of neural coreference resolution on the development set (Table 2). Our hypothesis is that D_{Lin} , the linear distance in elementary DUs, may not offer much more information than the linear distance in tokens that the neural model already employs. D_{LCA} , the distance from the right-hand mention to the LCA, on the other hand, may not be accurate when we artificially merge the RST trees for each paragraph into a single right-branched tree. In this case, the depth of the right-hand branch depends more on the order of paragraphs than the actual discourse structure of the text.

The mean results of the model enhanced with the rhetorical distances D_{Rh} are not much different from the baseline results on both sets. However, its results vary more, hence the model with the best F1 score reached both leaderboards. This suggests to us that the rhetorical distance is more robust than the other features, even though it shares all the mentioned drawbacks of the other features.

6 Conclusion

In this paper, we propose a new method for neural coreference resolution that incorporates discourse information. We test our method on the RuCoCo-2023 Shared Task and demonstrate that it outperforms the competition baseline by a significant margin, while also ranking 1st on the development set and 2nd on the test set of the competition. The key findings of this work are:

1. We implemented various features related to distances in the text-level RST tree to study how the hierarchical discourse information obtained with discourse parser can help coreference resolution for Russian.
2. We observed a marginal improvement using the rhetorical distance feature. The model that uses this feature got the best result on the Shared Task development and test sets.
3. We used the multilingual entity-aware LUKE model and showed that it performs competitively with the monolingual language models for Russian in coreference resolution, even with limited computational resources.

These findings suggest that the multilingual entity-aware LUKE model is a viable option for coreference resolution in Russian, and despite the constraints of the current rhetorical analyzer for Russian that prevent full-text analysis, the features of hierarchical discourse can still be found useful. We hope that our work will inspire further research on incorporating referential distance information into neural coreference resolution models.

Acknowledgements

The research was carried out using the infrastructure of the Shared Research Facilities “High Performance Computing and Big Data” (CKP “Informatics”) of FRC CSC RAS (Moscow). This study was conducted within the framework of the scientific program of the National Center for Physics and Mathematics, section №9 “Artificial intelligence and big data in technical, industrial, natural and social systems”.

References

- Elena Chistova, Artem Shelmanov, Dina Pisarevskaya, Maria Kobozeva, Vadim Isakov, Alexander Panchenko, Svetlana Toldova, and Ivan Smirnov. 2021. RST discourse parser for Russian: an experimental study of deep learning models. // *Analysis of Images, Social Networks and Texts: 9th International Conference, AIST 2020, Skolkovo, Moscow, Russia, October 15–16, 2020, Revised Selected Papers 9*, P 105–119. Springer.
- Alexis Conneau and Guillaume Lample. 2019. Cross-lingual language model pretraining. // *Proceedings of the 33rd International Conference on Neural Information Processing Systems*, Red Hook, NY, USA. Curran Associates Inc.
- Dan Cristea, Nancy Ide, Daniel Marcu, and Valentin Tablan. 1999. Discourse structure and co-reference: An empirical study. // *The Relation of Discourse/Dialogue Structure and Reference*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. // *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, P 4171–4186, Minneapolis, Minnesota, June. Association for Computational Linguistics.
- Vladimir Dobrovolskii, Mariia Michurina, and Alexandra Ivoylova. 2022. RuCoCo: a new Russian corpus with coreference annotation. // *COMPUTATIONAL LINGUISTICS AND INTELLECTUAL TECHNOLOGIES*. RSUH, June.
- Vladimir Dobrovolskii. 2021. Word-level coreference resolution. // *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, P 7670–7675, Online and Punta Cana, Dominican Republic, November. Association for Computational Linguistics.
- Olga Fedorova, Ekaterina Delikishkina, and Anna Uspenskaya. 2010. Experimental approach to reference in discourse: Working memory capacity and language comprehension in Russian. // *Proceedings of the 24th Pacific Asia Conference on Language, Information and Computation*, P 125–132, Tohoku University, Sendai, Japan, November. Institute of Digital Enhancement of Cognitive Processing, Waseda University.
- Talmy Givón. 1983. Topic continuity in discourse. *Topic continuity in discourse*, P 1–498.
- Barbara J. Grosz and Candace L. Sidner. 1986. Attention, intentions, and the structure of discourse. *Computational Linguistics*, 12(3):175–204.
- Sopan Khosla, James Fiacco, and Carolyn Rosé. 2021. Evaluating the impact of a hierarchical discourse representation on entity coreference resolution performance. // *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, P 1645–1651, Online, June. Association for Computational Linguistics.
- Andrej A Kibrik and Olga N Krasavina. 2005. A corpus study of referential choice: The role of rhetorical structure. *Computational Linguistics and Intellectual Technologies. Papers from the Annual International Conference “Dialogue” (2005)*, P 561–569.
- Andrej A Kibrik. 1999. Cognitive inferences from discourse observations: reference and working memory. // *Discourse studies in cognitive linguistics. Proceedings of the 5th International cognitive linguistics conference*, P 29–52.

- Yuval Kirstain, Ori Ram, and Omer Levy. 2021. Coreference resolution without span representations. // *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, P 14–19, Online, August. Association for Computational Linguistics.
- Kenton Lee, Luheng He, Mike Lewis, and Luke Zettlemoyer. 2017. End-to-end neural coreference resolution. // *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, P 188–197, Copenhagen, Denmark, September. Association for Computational Linguistics.
- Kenton Lee, Luheng He, and Luke Zettlemoyer. 2018. Higher-order coreference resolution with coarse-to-fine inference. // *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, P 687–692, New Orleans, Louisiana, June. Association for Computational Linguistics.
- Petter Mæhlum, Dag Haug, Tollef Jørgensen, Andre Kåsen, Anders Nøklestad, Egil Rønningstad, Per Erik Solberg, Erik Velldal, and Lilja Øvreliid. 2022. NARC – Norwegian anaphora resolution corpus. // *Proceedings of the Fifth Workshop on Computational Models of Reference, Anaphora and Coreference*, P 48–60, Gyeongju, Republic of Korea, October. Association for Computational Linguistics.
- 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.
- Nafise Sadat Moosavi and Michael Strube. 2016. Which coreference evaluation metric do you trust? a proposal for a link-based entity aware metric. // *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, P 632–642, Berlin, Germany, August. Association for Computational Linguistics.
- Megan Moser and Johanna D. Moore. 1996. Toward a synthesis of two accounts of discourse structure. *Computational Linguistics*, 22(3):409–419.
- Ryokan Ri, Ikuya Yamada, and Yoshimasa Tsuruoka. 2022. mLUKE: The power of entity representations in multilingual pretrained language models. // *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, P 7316–7330, Dublin, Ireland, May. Association for Computational Linguistics.
- Raghuveer Thirukovalluru, Nicholas Monath, Kumar Shridhar, Manzil Zaheer, Mrinmaya Sachan, and Andrew McCallum. 2021. Scaling within document coreference to long texts. // *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, P 3921–3931, Online, August. Association for Computational Linguistics.
- Liyan Xu and Jinho D. Choi. 2020. Revealing the myth of higher-order inference in coreference resolution. // *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, P 8527–8533, Online, November. Association for Computational Linguistics.
- Ikuya Yamada, Akari Asai, Hiroyuki Shindo, Hideaki Takeda, and Yuji Matsumoto. 2020. LUKE: Deep contextualized entity representations with entity-aware self-attention. // *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, P 6442–6454, Online, November. Association for Computational Linguistics.