CONSTITUENCY TREE REPRESENTATION FOR ARGUMENT UNIT RECOGNITION

Anonymous authors
Paper under double-blind review

ABSTRACT

The extraction of arguments from sentences is usually studied by considering only the neighbourhood dependencies of words. Such a representation does not rely on the syntactic structure of the sentence and can lead to poor results especially in languages where grammatical categories are scattered in the sentence. In this paper, we investigate the advantages of using a constituency tree representation of sentences for argument discourse unit (ADU) prediction. We demonstrate that the constituency structure is more powerful than simple linear dependencies between neighbouring words in the sentence. Our work was organised as follows: First, we compare the maximum depth allowed for our constituency trees. This first step allows us to choose an optimal maximum depth. Secondly, we combine this structure with graph neural networks, which are very successful in neural network tasks. Finally, we evaluate the benefits of adding a conditional random field to model global dependencies between labels, given local dependency rules. We improve the current best models for argument unit recognition at token level and also present an explainability method to evaluate the suitability of our model architecture.

1 INTRODUCTION

Argument Identification inside documents is the first step to study rhetorical processes in speeches or political debates. Recently, this domain has considerably gained attention due to ongoing progresses in various Natural Language Processing (NLP) tasks with the development of language models and transfer learning.

Due to the data overflow coming from social networks, argument analysis inside political debates is a concrete NLP application that interests the media community. The goal is to precisely identify arguments inside a sentence and predict its stance, given a topic.

Former works on argument analysis at the token level has generally used language models such as the Bidirectional Encoder Representations from Transformers model (BERT) (Devlin et al., 2019), associated to probabilistic models such as Conditional Random Fields (CRF). This combination enables to increase the global prediction coherence with restricted fine-tuning.

Language models are essential to incorporate statistical knowledge without having to train on large datasets. However, these models were recently criticized for their lack of transparency on how training data are reused in predictions. Thus, model architectures based on linguistic concepts seem to be needed to improve current models and increase their explainability. Despite various initiative to build bridges between linguists and data researchers (Gauthier et al., 2020), the most commonly used models do not do not take into account the grammatical structure of sentences, which leads to a gap between current NLP models and linguistic analysis.

In this work, we study the impact of incorporating grammatical structure into the model. We focus on the constituency tree representation of the sentence, as presented in Figure 1 as an example. This representation is composed of internal nodes representing the grammatical structure and leaf nodes (nodes without children) corresponding to the sentence words. This tree has been generated by BENEPAR (Kitaev & Klein, 2018) and trained by their authors on the SPMRL dataset. Note that we use such tree as a basis for our model without retraining.

The following paper mainly aims to:

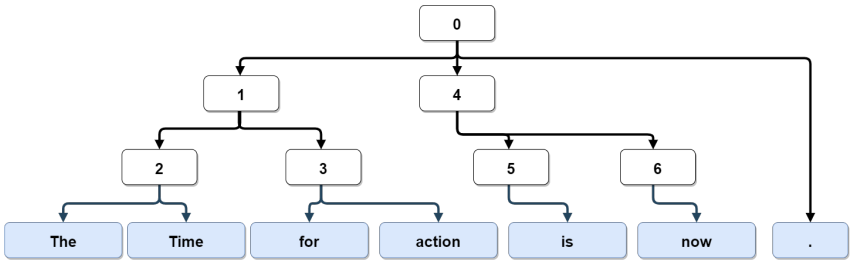

Figure 1: Constituency tree representation of the sentence "The time for action is now". The names of grammatical classes have been replaced by numbers (white nodes) and the depth of the tree was capped to a maximal value of 3.

- improve the current best model for identifying argumentative units by adding syntactic information from the sentence.
- evaluate the contribution of the constituency tree representation to the results.

2 Related Works

Argument Mining. Defining the concept of argument is an important step in creating dataset annotation rules. Previous work on Argument Mining (Lawrence & Reed, 2020) has identified two main scales of argument units: Elementary Discourse Units (EDUs), which are the non-overlapping text spans corresponding to the atomic discourse units, and Argumentative Discourse Units (ADUs), defined as minimal analysis units. As explained in Peldszus & Stede (2013), "when two EDUs are joined by some coherence relation that is irrelevant for argumentation, the resulting complex might be the better ADU". With this definition, one can notice that EDUs and ADUs are not defined at the sentence level. Therefore, it is more relevant to use a token-labeled dataset rather than a sentence-labeled dataset.

The current biggest dataset with argument unit ADUs labels at the token level is AURC [1] (Trautmann et al., 2020). They model arguments as self-contained pieces of information that can be verified as relevant arguments for a given topic with no or minimal surrounding context.

Tree Structure Representation in NLP. The goal of using a tree representation of the sentence is to exploit the syntactic hierarchical dependencies of the sentence with graph neural network. Thus, constituency tree representation is more appropriate than POS taggers which provide one tag per token (according to the Universal POS tags categories).

Pioneer work (Gildea & Palmer, 2002) presenting the benefits of using constituency tree representation has failed to scale into production. This is caused by the lack of a reliable model to generate constituency tree representation of the sentences. However, recent promising results were made in consistency and dependency parsing.

For our preprocessing step we used a neural network model called Berkeley Neural Parser (BENEPAR) (Kitaev & Klein, 2018), which is trained on 11 languages. We use this model with the weights provided by these designers. Other possibilities for this task are the StanfordParser with Stanza Qi et al. (2020) or NLTK . We chose BENEPAR because it is integrated with Spacy and it can be used in 11 different languages included English and French, which are our main language of interest for our research.

Other papers have recently studied the use of tree structure to incorporate syntactic information to their models. Marcheggiani & Titov (2017) uses Graph Convolutional Networks (GCNs) based on the dependency tree structure of the sentence for semantic role labeling. Beck et al. (2018) uses GNNs for generation tasks from abstract meaning representation.

Graph Neural Network. Graph Neural Networks refer to the various methods which extend Deep Learning to data organized in graphs. GNNs were introduced by Gori et al. (2005) as an extension

[1] The dataset is publicly available at `https://github.com/trtm/AURC`

of Recursive Neural Networks (RNN). GNNs use a neural network to propagate information within the dataset geometrical structure. In this article, the Graph Attention Layer (GAT) (Veličković et al., 2018) is used. It appears to be the most adapted message passing layer to propagate information following a self-attention strategy.

Conditional Random Field. Conditional Random Fields were introduced by (Hammersley & Clifford, 1971) as probabilistic models for magnetic field modeling. Their goal was to exploit the Markov assumption of a lattice to derive the log probability of the label distribution formula. Since then, CRFs have been widely used in machine learning thanks to the iterative scaling algorithms presented by Lafferty et al. (2001) to find the parameter that maximizes the log-likelihood of the training data. Many interactions have been made between CRFs and GNNs. Gao et al. (2019) proposed a new CRF layer for GCNs. Cui & Zhang (2019) proposed a neural network alternative to CRFs motivated by the fact that the Markov assumption is not always respected for linear chain representation of the sentence.

3 Experimental Setup

3.1 Data Source and Annotation Rules

The dataset used in our study has been introduced by Trautmann et al. (2020). This dataset appears to be one of the few datasets available that annotate argument unit at the token level. It is composed of 8000 sentences, equally divided into 8 topics, and can be used for the task of argument unit identification as well as argument polarity identification.

The authors distinguished between PRO (supporting), CON (opposing) arguments and NON (non-argumentative) words. The data are also annotated at the sentence level according to the following rules:

- if all the words have the same label, the sentence also has this label;
- if some words have the labels NON and the others have all the same label PRO (resp. CON), the sentence has the label PRO (resp. CON);
- if some words of the sentence have the label PRO and others have the label CON, the sentence has the majority label (the label is chosen randomly between PRO and CON if proportions are equal).

This implies an asymmetry between the NON label and the other two labels: the PRO and CON labels identify arguments and are predominant over the NON label which indicates the absence of an argument.

In our pre-processing step, we needed to label the internal nodes of the constituency tree in order to use a graphical neural network that predicts the labels of the nodes. As these annotations are not provided by the dataset, we chose to annotate the internal nodes with the same labelling rules as described above for sentence labelling. This preserves the internal logic of the labelling process.

3.2 Constituency tree construction

One of the main advantages of adding a constituency tree to argument identification methods is the greater proximity of words that belong to the same grammatical class. Indeed, we noticed that words belonging to the same grammatical class have more often the same label than words that are only neighbours in the sense of the linear representation of the sentence.

To evaluate this assumption, we calculated three proportions to better assess the proximity of nodes within a constituency tree. The three calculated proportions, summarised in Table 1, are as follows.

- **Neighbour similarity**: the proportion of nodes which have the same label and are neighbours in the linear chain sentence representation. For the cases (columns 2, 3, 4) where we have also a constituency tree representation, we additionally restrict this set to the nodes which share the same grammatical class (i.e. which have the same parent node).
- **Similarity of grammatical classes and words**: the proportion of words which have the same label as their grammatical class (the parent node in the sens of a tree).

- **Similarity of grammatical classes**: the proportion of internal nodes which are linked by an edge and share the same label.

We present the results according to the maximum depth allowed for the constituency tree. When a tree is deeper than the maximum depth allowed, we have capped the tree and all nodes in the capped branch are neighbours.

Category \ Maximal depth of the tree	No tree	2	3	4
Neighbour similarity	96.8 %	98.2%	98.3%	**98.5**%
Similarity of grammatical classes and words	//	89%	**90**%	83%
Similarity of grammatical classes	//	88%	**93**%	92%

Table 1: Evaluation of three proportions to assess the appropriateness of using the constituency tree representation.

Table 1 shows a stronger proximity between neighbour words in the same grammatical class than neighbour words when the grammatical structure is not taken into account. We can also notice that the constituency tree with a maximal depth of 4 shows a stronger similarity of neighbouring node than the two others. However, 4 is not necessarily the best depth as the constituency tree with a maximal depth of 3 shows better results concerning the grammatical class similarity. This lead us to try our models with the maximum depth of 3 and 4.

This fact is also illustrated with the constituency tree shown in Figure 2. Within this sentence, we can see that the ADU units are distributed according to the grammatical structure of the sentence. The words "inflation" and "the" are neighbours in the sentence but are further apart in the constituency tree. This leads to a better identification of the boundaries between the ADUs.

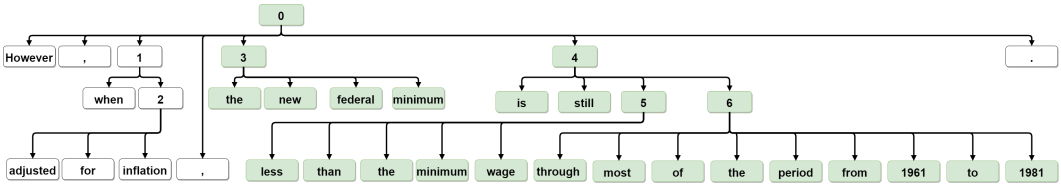

Figure 2: Constituency tree representation of the sentence : *"However, when adjusted for inflation, the new federal minimum is still less than the minimum wage through most of the period from 1961 to 1981."* with a maximum depth allowed of 3. The green nodes represents words or spans with "PRO" label and the white nodes represents words and spans with "NON" label.

For the construction of the constituency tree, we use Berkeley Neural Parser (BENEPAR) (Kitaev & Klein, 2018). This is a multilingual constituency parser which benefits from unsupervised pre-training across a variety of languages. This model provides the weights for 11 languages and a simple API incorporated to Spacy, and is thus a good choice in order to have a good accuracy, integrated into a larger pipeline. This additional pre-processing step requires to have an exact match between the tokenizations used by BENEPAR and BERT models. This constraint imposed us to restrict the dataset from 8000 to 6709 sentences (the size of the development dataset is now 790, the size of the training dataset is now 3960, the size of the test dataset is now 1959).

3.3 Evaluation Score

In this study, we used the F1-score as a metric for evaluating the models. The F1-score corresponds to the equally weighted harmonic mean of the precision and recall. There are two ways to calculate the F1 score. The micro-average calculates the metric globally by counting the total of true positives, false negatives and false positives. The macro-average calculates the metrics for each label and finds their unweighted average. It does not take into account label imbalance. Therefore, the micro-average is preferable for our studies, as the database is unbalanced at the class level with respect

to the non-argumentative label at the token level. The F1-score is calculated at the token level of BERT. As some words are split into two tokens with the BERT tokenizer, the precision and recall ratings are slightly biased. One way to minimise this effect is to merge the labels of the tokens of a word that has been split.

4 MODEL ARCHITECTURE

4.1 PREVIOUS MODEL

The best previous model for identifying the argumentative unit at the token level was proposed by Trautmann et al. (2020).

It consists of two steps. In the first step, the sentence is tokenized by following the BERT tokenizer where the BERT model is fitted for token classification. The result is equal in size to the number of label classes. In a second step, a linear chain conditional random field is applied to improve the prediction by taking into account the label of these neighbours.

The main idea of this model is to take advantage of the knowledge of BERT's attention and then to improve the results by harmonising the labels of neighbouring words. This takes advantage of neighbours dependency relations between words.

4.2 BERT-GAT-CRF

In order to take advantage of the dependency structure of the sentence, we incorporate a GNN between the use of BERT and the linear chain CRF. The idea behind this architecture is as follows.

First, the BERT language model produces the hidden representation of the sentence. Then, the information is disseminated to the neighbors of the graph at each iteration. In this way, we hope to achieve better consistency between neighboring nodes when training on a small dataset. A major difficulty in choosing a graph neural layer is that each sentence has a different tree representation. The message-passing design allows the model weights to be shared among the nodes in the network, so that the results do not depend on the initial access to the global tree structure. It also allows for both directed and undirected graphs, as the flow of information is conditional on the existence of a directed edge between two nodes.

The attention mechanism was introduced by Vaswani et al. (2017) for translation tasks in NLP and has since been widely used in many other tasks. The main advantages of attention mechanisms are that they can handle inputs of varying size and specify different weights for different neighbors by focusing on the most relevant parts of the input to make decisions.

The Graph Attention Layer (GAT) (Veličković et al., 2018) allows the attention mechanism to be combined with the graph structure while preserving the syntactic structural information of the sentence. In order to improve the stability of the model, the addition of multi-head attention layers could be beneficial in the training stage. The different heads are then aggregated to provide the next hidden states of the neural network.

Unlike the graph convolution layer, this layer allows nodes in the same neighborhood to be assigned different importance. In addition, the analysis of the learned attention weights can improve the interpretability of the model, as we can extract information from the attention weight matrices.

Note that we are aware of the potential for improvement by replacing BERT with ROBERTA Liu et al. (2019) or another language model that has outperformed BERT. We chose to use BERT as our language model so that we could compare our results with the previous best model regarding only the addition of syntactic structure. Moreover, we used BERT from the Transformer library developed by Hugging Face Wolf et al. (2020). It is thus easy to replace BERT by another language model without reimplementing the whole network.

4.3 HYPERPARAMETERS OPTIMIZATION

The BERT-GAT-CRF model have many more hyperparameters than the BERT-CRF model. This is due to the large amount of hyperparameters of a GAT network: number of layer, units per layer, number of heads.

In order to choose the best hyperparameters for this model, we use the library Optuna (Akiba et al., 2019). Optuna is a framework for cost-effectiveness of hyperparameters optimization. The two main advantages of Optuna for our study are the following.

- Searching strategy: Optuna implements a relational sampling method that can identify the underlying concurrence relations after some independent samplings.
- Pruning strategy: regarding the intermediate F1 score values, Optuna interrupts the non promising trials based on the Successive Halving algorithm Li et al. (2020), which continues the trials if only the actual F1-scores is among the best intermediate results.

The importance of the different hyperparameters of our model has been evaluated in Table 2. It can be observed that the most important hyperparameters are the learning rate and the maximum allowed value for the gradient. From an empirical point of view, we notice that when we do not constrain the gradient, the model converges to a local optimum which is the assignment of the label NO to each word. This local optimum exists because the dataset is unbalanced towards the absence of an argument.

	Test Intervals	Best values	Parameters importance
Learning rate	10^{-5} to 10^{-3}	$2.8 \cdot 10^{-5}$	30 %
Maximum gradient allowed	10^{-1} to 10^{2}	9.7	49 %
Number of GAT layers	1 to 3	2	2 %
Number of unit per GAT layers	50 to 300	290 and 100	2 %
Number of heads per GAT layers	1 to 3	3 and 3	7 %
Number of linear layers	1 to 3	2	5 %
Number of unit per linear layers	50 to 250	100 and 100	5 %

Table 2: Feature importance of the BERT-GAT-CRF model in Cross Domain prediction

The main difficulty encountered during the development of this model was the management of batches. In Pytorch Geometric, we associate a graph with each input. This graph is represented by a two-dimensional matrix that contains the edges between the nodes. Thus, in order to allow different sizes of graphs, the batches matrix used in Pytorch is unfolded into a list and a list of indices links the input data to the batch indices. Thus, as we have to switch between the two representations in order to use BERT and a GNNs within our model, this reduces the computational speed.

The global training step requires loading the BERT model and a batch of data into memory. According to our implementation, this requires 30 GB of memory (GPU or CPU). We therefore performed our global training experiments on the CPU, which took about 2 hours per iteration and led to about 30 iterations per model configuration. The transfer learning model training required about 7 GB and could therefore be run on the GPU. We therefore had about 200 trials per model hyper-parameter optimisation.

5 MODELS RESULTS

5.1 MODELS EVALUATION

Consider first the performance of the models when the BERT model is not fine tuned. In this case, the model BERT-CRF shows poor result: 31% of F1 score probably due to the low number of parameters (we only train here the last layer of BERT and the parameters of the CRF). The BERT-GNN-CRF model reaches 50% of F1-score on the Cross Domain dataset.

In a second step, we perform a global training of the models, whose results are presented in Table 3. As suggested by Table 1, we compute our models with a maximum depth allowed for the constituency tree of 3 and 4. We reach our best results with the model composed of BERT-GNN-CRF

	In Domain	Cross Domain
BERT large	68 % *	59 % *
BERT large - Linear chain CRF	69 % *	61 % *
BERT large - GAT - CRF - Maximum depth allowed for Constituency tree is 3	72.8 %	68.2 %
BERT large - GAT - CRF - Maximum depth allowed for Constituency tree is 4	**73.5** %	**68.7** %

Table 3: F1- score of the different models at token level on the test dataset (the results with a star (*) have been directly taken from (Trautmann et al., 2020))

when we allow a maximal depth of 4, which achieves 74.1 % on token level F1-score micro in a cross-domain setting and 73.8 % on token level F1 score micro for the in-domain setup.

5.2 Model Interpretability

In this subsection, we study the interpretability of the model BERT-GNN-CRF.

Consider the problem of attributing the model prediction to its input features. First, we use the Integrated Gradient method (Sundararajan et al., 2017) which computes a score for every input of the model compared to a baseline value. The score is calculated by integrating the gradient along the input representation from the baseline to our input. When this gradient is higher, our input data is more relevant for prediction. In the BERT model, this baseline is the token "[PAD]" (commonly used for padding) which indicates the absence of a word.

We used the Captum framework (Kokhlikyan et al., 2020) which implements a wide variety of interpretability models. Figure 3 presents the contribution of the words of the sentence to the final output. In this example, it can be seen that the word contributions follow the grammatical structure (expect for the word "then", due to the fact that labeling did not follow strictly the grammar).

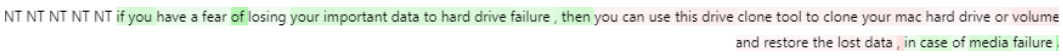

Figure 3: Word Importance of the Integrated Gradient Method for the sentence : *"If you have a fear of losing your important data to hard drive failure, then you can use this drive clone tool to clone your mac hard drive or volume and restore the lost data, in case of media failure."*. The "NT" token refers to the sentence internal nodes. The words highlighted in green (resp. red) have a positive (resp. negative) impact on the final outputs.

Then, we study the importance of the edges for the model prediction. The Integrated Gradient method cannot be directly applied to our graph study as we cannot integrate through the edges weights (the edges are not weighted). We compute a simple algorithm to measure the importance of the different edges of our graph. We iteratively hide one edge from the graph and compute the model to see how the result varies. By doing so with all the edges we can measure the necessity of every edges taken independently. The algorithm is described in the Algorithm 1.

Algorithm 1 Measure of the impact of the graph edges.

Input: A sentence to analyse, The edges of the constituency tree of the sentence, The labels of the sentence
Output: matrix containing the importance weight of each edge

```
Run the model on the input sentence
for Every edges of the graph do
    Remove one edge of the graph
    Run the model with the new input
    Compute difference between the previous loss and the new loss
    Store this value as the importance of the edge
end for
```

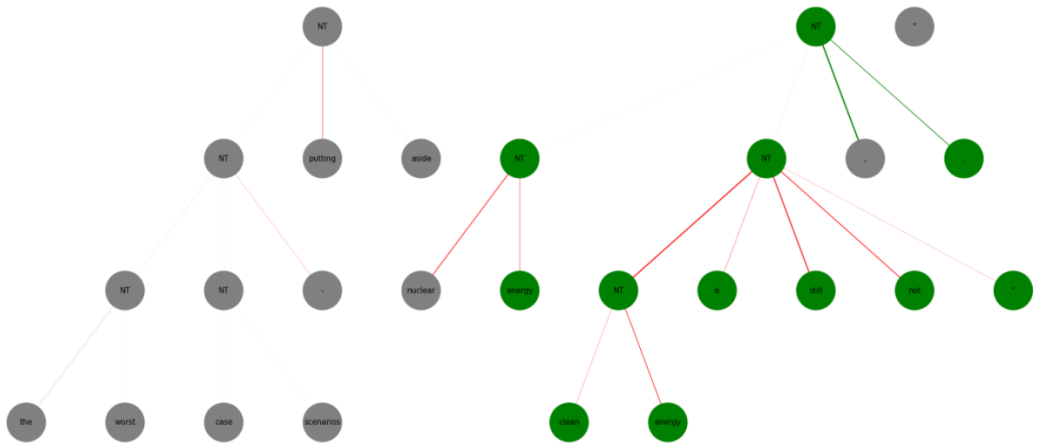

Figure 4: Tree representation of the edges interpretability methods. The red edges refer to edges which have a positive impact on the prediction. The green edges refer to edges which have a negative impact on the prediction. (This figure has been generated by NetworkX)

The results of the interpretability on edges can be found in Figure 4. It globally shows that the words which are labeled with an argument (in color) furnish more information to their related edges. This confirm the utility of adding the constituency tree representation to our model.

6 Conclusion

In this study, we established a new method for identifying the ADU boundary using the sentence constituent tree representation. This model allows the information to be spread across the graph and performs well on a small dataset.

The two main errors previously observed on this dataset were:

- the span of an argumentative segments was not correctly recognized.
- the stances are not correctly classified.

In this paper, we focus on the span detection problem and improve the method for identifying the boundaries of an ADU. However, we do not have a strict adherence to grammatical correctness. As mentioned in Trautmann et al. (2020), this would require that all spans containing arguments be clauses. There are examples of sentences where subjects are absent, which limits the performance of models using grammatical structure. This problem could be solved by specifying annotation rules that more strictly follow the syntactic structure of the sentence.

Regarding the second type of error, the problem of position identification is mainly related to the performance of the BERT model. The dataset only provides sentences with a maximum length of 64. This implies that only a narrow context of the ADU is available, which limits the capability of our model. Many arguments require more domain knowledge to understand what is at stake.

They model arguments as self-contained pieces of information that can be verified as relevant arguments for a given topic with no or minimal surrounding context.

A disadvantage of using message-passing GNNs is that, since not all words have the same depth in tree representations, message passing is treated identically across words. A word with shallow depth will pass information to a much more distant part of the tree than a node with greater depth.

In future work, we will explore other GNNs to overcome this problem. We plan to continue our study on argument recognition by working on datasets with more surrounding context.

REFERENCES

Takuya Akiba, Shotaro Sano, Toshihiko Yanase, Takeru Ohta, and Masanori Koyama. Optuna: A Next-generation Hyperparameter Optimization Framework. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, KDD '19, pp. 2623–2631, New York, NY, USA, July 2019. Association for Computing Machinery. ISBN 978-1-4503-6201-6. doi: 10.1145/3292500.3330701. URL https://doi.org/10.1145/3292500.3330701.

Daniel Beck, Gholamreza Haffari, and Trevor Cohn. Graph-to-Sequence Learning using Gated Graph Neural Networks. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 273–283, Melbourne, Australia, July 2018. Association for Computational Linguistics. doi: 10.18653/v1/P18-1026. URL https://www.aclweb.org/anthology/P18-1026.

Leyang Cui and Yue Zhang. Hierarchically-Refined Label Attention Network for Sequence Labeling. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pp. 4115–4128, Hong Kong, China, November 2019. Association for Computational Linguistics. doi: 10.18653/v1/D19-1422. URL https://aclanthology.org/D19-1422.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *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)*, pp. 4171–4186, Minneapolis, Minnesota, June 2019. Association for Computational Linguistics. doi: 10.18653/v1/N19-1423. URL https://aclanthology.org/N19-1423.

Hongchang Gao, Jian Pei, and Heng Huang. Conditional Random Field Enhanced Graph Convolutional Neural Networks. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, KDD '19, pp. 276–284, New York, NY, USA, July 2019. Association for Computing Machinery. ISBN 978-1-4503-6201-6. doi: 10.1145/3292500.3330888. URL https://doi.org/10.1145/3292500.3330888.

Jon Gauthier, Jennifer Hu, Ethan Wilcox, Peng Qian, and Roger Levy. SyntaxGym: An Online Platform for Targeted Evaluation of Language Models. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pp. 70–76, Online, July 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.acl-demos.10. URL https://aclanthology.org/2020.acl-demos.10.

Daniel Gildea and Martha Palmer. The Necessity of Parsing for Predicate Argument Recognition. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pp. 239–246, Philadelphia, Pennsylvania, USA, July 2002. Association for Computational Linguistics. doi: 10.3115/1073083.1073124. URL https://aclanthology.org/P02-1031.

M. Gori, G. Monfardini, and F. Scarselli. A new model for learning in graph domains. In *Proceedings. 2005 IEEE International Joint Conference on Neural Networks, 2005.*, volume 2, pp. 729–734 vol. 2, July 2005. doi: 10.1109/IJCNN.2005.1555942. ISSN: 2161-4407.

J.M Hammersley and P Clifford. Markov Fields On Finite Graphs And Lattices. 1971. URL http://www.statslab.cam.ac.uk/~grg/books/hammfest/hamm-cliff.pdf.

Nikita Kitaev and Dan Klein. Constituency Parsing with a Self-Attentive Encoder. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 2676–2686, Melbourne, Australia, July 2018. Association for Computational Linguistics. doi: 10.18653/v1/P18-1249. URL https://aclanthology.org/P18-1249.

Narine Kokhlikyan, Vivek Miglani, Miguel Martin, Edward Wang, Bilal Alsallakh, Jonathan Reynolds, Alexander Melnikov, Natalia Kliushkina, Carlos Araya, Siqi Yan, and Orion Reblitz-Richardson. Captum: A unified and generic model interpretability library for PyTorch. *arXiv:2009.07896 [cs, stat]*, September 2020. URL http://arxiv.org/abs/2009.07896. arXiv: 2009.07896.

John D. Lafferty, Andrew McCallum, and Fernando C. N. Pereira. Conditional Random Fields: Probabilistic Models for Segmenting and Labeling Sequence Data. In *Proceedings of the Eighteenth International Conference on Machine Learning*, ICML '01, pp. 282–289, San Francisco, CA, USA, June 2001. Morgan Kaufmann Publishers Inc. ISBN 978-1-55860-778-1.

John Lawrence and Chris Reed. Argument Mining: A Survey. *Computational Linguistics*, 45 (4):765–818, January 2020. ISSN 0891-2017, 1530-9312. doi: 10.1162/coli_a_00364. URL https://www.mitpressjournals.org/doi/abs/10.1162/coli_a_00364.

Liam Li, Kevin Jamieson, Afshin Rostamizadeh, Ekaterina Gonina, Jonathan Ben-tzur, Moritz Hardt, Benjamin Recht, and Ameet Talwalkar. A System for Massively Parallel Hyperparameter Tuning. *Proceedings of Machine Learning and Systems*, 2:230–246, March 2020. URL https://proceedings.mlsys.org/paper/2020/hash/f4b9ec30ad9f68f89b29639786cb62ef-Abstract.html.

Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. RoBERTa: A Robustly Optimized BERT Pretraining Approach. *arXiv:1907.11692 [cs]*, July 2019. URL http://arxiv.org/abs/1907.11692. arXiv: 1907.11692.

Diego Marcheggiani and Ivan Titov. Encoding Sentences with Graph Convolutional Networks for Semantic Role Labeling. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pp. 1506–1515, Copenhagen, Denmark, September 2017. Association for Computational Linguistics. doi: 10.18653/v1/D17-1159. URL https://www.aclweb.org/anthology/D17-1159.

Andreas Peldszus and Manfred Stede. From Argument Diagrams to Argumentation Mining in Texts: A Survey. *International Journal of Cognitive Informatics and Natural Intelligence*, 7(1):1–31, January 2013. ISSN 1557-3958. doi: 10.4018/jcini.2013010101. URL https://doi.org/10.4018/jcini.2013010101.

Peng Qi, Yuhao Zhang, Yuhui Zhang, Jason Bolton, and Christopher D. Manning. Stanza: A Python Natural Language Processing Toolkit for Many Human Languages. *arXiv:2003.07082 [cs]*, April 2020. URL http://arxiv.org/abs/2003.07082. arXiv: 2003.07082.

Mukund Sundararajan, Ankur Taly, and Qiqi Yan. Axiomatic Attribution for Deep Networks. In *International Conference on Machine Learning*, pp. 3319–3328. PMLR, July 2017. URL http://proceedings.mlr.press/v70/sundararajan17a.html. ISSN: 2640-3498.

Dietrich Trautmann, Johannes Daxenberger, Christian Stab, Hinrich Schütze, and Iryna Gurevych. Fine-Grained Argument Unit Recognition and Classification. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(05):9048–9056, April 2020. ISSN 2374-3468. doi: 10.1609/aaai.v34i05.6438. URL https://ojs.aaai.org/index.php/AAAI/article/view/6438. Number: 05.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is All you Need. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (eds.), *Advances in Neural Information Processing Systems 30*, pp. 5998–6008. Curran Associates, Inc., 2017. URL http://papers.nips.cc/paper/7181-attention-is-all-you-need.pdf.

Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. Graph Attention Networks. February 2018. URL https://openreview.net/forum?id=rJXMpikCZ.

Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. Transformers: State-of-the-Art Natural Language Processing. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pp. 38–45, Online, October 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.emnlp-demos.6. URL https://aclanthology.org/2020.emnlp-demos.6.