

Look Again at the Syntax: Relational Graph Convolutional Network for Gendered Ambiguous Pronoun Resolution

Yinchuan Xu*

University of Pennsylvania
Philadelphia, PA, 19104, USA
yinchuan@seas.upenn.edu

Junlin Yang*

Yale University
New Haven, CT, 06511, USA
junlin.yang@yale.edu

Abstract

Gender bias has been found in existing coreference resolvers. In order to eliminate gender bias, a gender-balanced dataset Gendered Ambiguous Pronouns (GAP) has been released and the best baseline model achieves only 66.9% F1. Bidirectional Encoder Representations from Transformers (BERT) has broken several NLP task records and can be used on GAP dataset. However, fine-tune BERT on a specific task is computationally expensive. In this paper, we propose an end-to-end resolver by combining pre-trained BERT with Relational Graph Convolutional Network (R-GCN). R-GCN is used for digesting structural syntactic information and learning better task-specific embeddings. Empirical results demonstrate that, under explicit syntactic supervision and without the need to fine tune BERT, R-GCN’s embeddings outperform the original BERT embeddings on the coreference task. Our work significantly improves the snippet-context baseline F1 score on GAP dataset from 66.9% to 80.3%. We participated in the Gender Bias for Natural Language Processing 2019 shared task, and our codes are available online.¹

1 Introduction

Coreference resolution aims to find the linguistic mentions that refer to the same real-world entity in natural language (Pradhan et al., 2012). Ambiguous gendered pronoun resolution is a subtask of coreference resolution, where we try to resolve gendered ambiguous pronouns in English such as "he" and "she". This is an important task for natural language understanding and a longstanding challenge. According to (Sukthanker et al., 2018), there are two main approaches: heuristics-based

approaches and learning-based approaches, such as mention-pair models, mention-ranking models, and clustering models (McCarthy and Lehnert, 1995; Haghighi and Klein, 2010; Fernandes et al., 2014). Learning-based approaches, especially deep-learning-based methods, have shown significant improvement over heuristics-based approaches.

However, most state-of-art deep-learning-based resolvers utilize one-directional Transformers (Stojanovski and Fraser, 2018), limiting the ability to handle long-range inferences and the use of cataphors. Bidirectional Encoder Representations from Transformers, or BERT (Devlin et al., 2018) learns a bidirectional contextual embedding and has the potential to overcome these problems using both the previous and next context. However, fine-tuning BERT for a specific task is computationally expensive and time-consuming.

Syntax information has always been a strong tool for semantic tasks. Most heuristics-based methods use syntax-based rules (Hobbs, 1978; Lappin and Leass, 1994; Haghighi and Klein, 2009). Many of learning based models also rely on syntactic parsing for mention or entity extraction algorithms and compute hand-crafted features as input (Sukthanker et al., 2018).

Can we learn better word embeddings than BERT on the coreference task with the help of syntactic information and without computationally expensive fine-tuning of BERT? Marcheggiani and Titov et al. (2017) successfully use Graph Convolutional Networks (GCNs) (Duvenaud et al., 2015; Kipf and Welling, 2016) to learn word embeddings for the semantic role labeling task and outperform the original LSTM contextual embeddings.

Inspired by Marcheggiani and Titov (2017), we create a 'Look-again' mechanism which combines BERT with Gated Relational Graph Convolutional

* Equal contribution.

¹Our codes and models are available at: <https://github.com/ianycxu/RGCN-with-BERT>.

Networks (R-GCN) by using BERT embeddings as initial hidden states of vertices in R-GCN. R-GCN’s structure is derived from a sentence’s syntactic dependencies graph. This architecture allows contextual embeddings to be further learned and encoded into better task-specific embeddings without fine tuning BERT which is computationally expensive.

2 Contributions

Our main contributions are: (1) Our work is the first successful attempt of using R-GCN to boost the performance of BERT contextual embeddings without the need to fine tune BERT. (2) Our work is the first to use R-GCN on the coreference resolution task. (3) Our work improves the snippet-context baseline F1 score on Gendered Ambiguous Pronouns dataset from 66.9% to 80.3%.

3 Methodology

We propose a series connection architecture of pre-trained BERT with Gated Relational Graph Convolutional Network (Gated R-GCN). Gated R-GCN is used for digesting structural syntactic information. This architecture, which we name as ‘Look-again’ mechanism can help us learn embeddings which have better performance on coreference task than original BERT embeddings.

3.1 Syntactic Structure Prior

As mentioned in the Introduction section, syntactic information is beneficial to semantic tasks. However, how to encode syntactic information directly into deep learning systems is difficult.

Marcheggiani and Titov (2017) introduces a way of incorporating syntactic information into sequential neural networks by using GCN. The syntax prior is transferred into a syntactic dependency graph, and GCN is used to digest this graph information. This kind of architecture is utilized to incorporate syntactic structure prior with BERT embeddings for coreference task in our work.

3.2 GCN

Graph Convolutional Networks (GCNs) (Duvenaud et al., 2015; Kipf and Welling, 2016) take graphs as inputs and conduct convolution on each node over their local graph neighborhoods. The convolution process can also be regarded as a simple differentiable message-passing process. The message here is the hidden state of each node.

Consider a directed graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ with nodes $v_i \in \mathcal{V}$ and edges $(v_i, v_j) \in \mathcal{E}$. The original work of GCN (Kipf and Welling, 2016) assumes that every node v contains a self-loop edge, which is $(v_i, v_i) \in \mathcal{E}$. We denote hidden state or features of each node v_i as h_i , and neighbors of each node as $\mathcal{N}(v_i)$, then for each node v_i , the feed-forward processing or message-passing processing then can be written as:

$$h_i^{(l+1)} = ReLU \left(\sum_{u \in \mathcal{N}(v_i)} \frac{1}{c_i} W^{(l)} h_u^{(l)} \right) \quad (1)$$

Note that we ignore the bias term here. l here denotes the layer number, and c_i is a normalization constant. We use $c_i = |\mathcal{N}(v_i)|$, which is the in-degree of the node. Weight $W^{(l)}$ is shared by all edges in layer l .

3.3 R-GCN

Each sentence is parsed into its syntactic dependencies graph and use GCN to digest this structural information. Mentioned in (Schlichtkrull et al., 2018), when we construct the syntactic graph we also allow the information to flow in the opposite direction of syntactic dependency arcs, which is from dependents to heads. Therefore, we have three types of edges: first, from heads to dependents; second, from dependents to heads and third, self-loop (see Fig. 1).

Traditional GCN cannot handle this multi-relation graph. Schlichtkrull (2018) proposed a Relational Graph Convolutional Networks (R-GCNs) structure to solve this multi-relation problem:

$$h_i^{(l+1)} = ReLU \left(\sum_{r \in R} \sum_{u \in \mathcal{N}_r(v_i)} \frac{1}{c_{i,r}} W_r^{(l)} h_u^{(l)} \right) \quad (2)$$

where $\mathcal{N}_r(v_i)$ and $W_r^{(l)}$ denote the set of neighbor of node i and weight under relation $r \in R$ respectively. In our case, we have three relations.

3.4 Gate Mechanism

Because the syntax information is predicted by some NLP packages, which might have some error, we need some mechanism to reduce the effect of erroneous dependency edges.

A gate mechanism is introduced in (Marcheggiani and Titov, 2017; Dauphin et al., 2017; Li




Figure 1: Syntactic dependencies graph with three relations

et al., 2015). The idea is calculating a gate value ranging from 0 to 1, and multiplying it with the incoming message. The gate value is computed by:

$$g_{u,v}^{(l)} = \text{Sigmoid} \left(h_u^{(l)} \cdot W_{r,g} \right) \quad (3)$$

The final forward process of Gated R-GCN is:

$$h_i^{(l+1)} = \text{ReLU} \left(\sum_{r \in R} \sum_{u \in \mathcal{N}_r(v_i)} g_{u,v_i}^{(l)} \frac{1}{c_{i,r}} W_r^{(l)} h_u^{(l)} \right) \quad (4)$$

3.5 Connect BERT and R-GCN in Series

We use pre-trained BERT embeddings (Devlin et al., 2018) as our initial hidden states of vertices in R-GCN. This series connection between pre-trained BERT and Gated R-GCN forms the 'Look-again' mechanism. After pre-trained BERT encodes tokens' embeddings, Gated R-GCN will 'look again' at the syntactic information which is presented as structural information and further learn semantic task-specific embeddings with the explicit syntactic supervision by syntactic structure.

A fully-connected layer in parallel with Gated R-GCN is utilized to learn a compact representation of BERT embeddings of two mentions (A and B) and the pronoun. This representation is then concatenated with Gated R-GCN's final hidden states of those three tokens. The reason of concatenating R-GCN's hidden states with BERT embeddings' compact representation is that graph convolution of the GCN model is actually a special form of Laplacian smoothing (Li et al., 2018), which might mix the features of vertices and make them less distinguishable. By concatenation, we maintain some original embeddings information. After concatenation, we use a fully-connect layer for the final prediction. The visualization of the final end-to-end model is shown in Fig. 2.

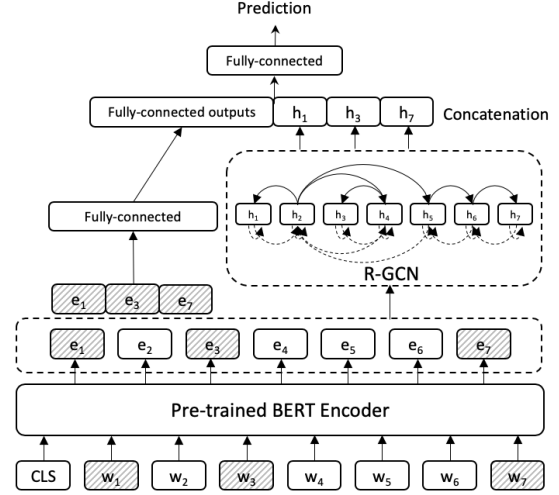


Figure 2: End-to-end coreference resolver

4 Experimental Methodology and Results

In the experiment, it shows that, with the explicit syntactic supervision by syntactic structure, Gated R-GCN structure can learn better embeddings that improve performance on the coreference resolution task. Two sets of experiments were designed and conducted: Stage one experiments and Full GAP experiments.

Stage one experiments used the same setting as stage one of shared-task competition, where we had 4454 data samples in total. 'Gap-validation.tsv' and 'gap-test.tsv' were used as training dataset, while 'gap-development.tsv' was used for testing.²

Full GAP experiments used full 8908 samples of Gendered Ambiguous Pronouns (GAP) dataset in order to compare with the baseline result from the GAP paper (Webster et al., 2018).

4.1 Dataset

The dataset provided by the shared task is Google AI Language's Gendered Ambiguous Pronouns (GAP) dataset (Webster et al., 2018), which is a gender-balanced dataset containing 8,908 coreference-labeled pairs of (ambiguous pronoun, antecedent name), sampled from Wikipedia.

In stage one of the shared task, only 2454 samples were used as the training dataset, and 2000 samples were used as the test dataset.

²<https://www.kaggle.com/c/gendered-pronoun-resolution/data>.

4.2 Data Preprocessing

SpaCy was used as our syntactic dependency parser. Deep Graph Library (DGL)³ was used to transfer each dependency graph into a DGL graph object. Several graphs were grouped together as a larger DGL batch-graph object for batch training setting. R-GCN model was also implemented with DGL.

4.3 Training settings

Adam was used (Kingma and Ba, 2014) as our optimizer. Learning rate decay was applied. $l2$ regularization of both R-GCN’s and fully-connected layer’s weights was added to the training loss function. Batch-normalization and drop-out were used in all fully-connection layers. We used one layer for R-GCN which captures immediate syntactic neighbors’ information. BERT in our model was not fined tuned and was fixed for training. We used ‘bert-large-uncased’ version of BERT for generating original embeddings.

The five-fold ensemble was used to achieve better generalization performance and more accurate estimation of the model’s performance. The training dataset was divided into 5 folds. Each time of training we trained our model on 4 folds and chose the model which had the best validation performance on the left fold. This best model then was used to predict the test dataset. In the end, predicted results from 5 folds were averaged as the final result.

4.4 Stage One Experiments

There are 4 different settings for Stage One experiments for comparisons (see Fig. 3):

1. Only BERT embeddings are fed into an additional MLP for prediction.
2. Connect BERT with Gated R-GCN, but only feed Gated R-GCN’s hidden states into MLP for prediction.
3. Connect BERT with R-GCN, and the concatenation is fed into MLP for prediction. The gate mechanism is not applied to R-GCN
4. Connect BERT with Gated R-GCN, and the concatenation is fed into MLP for prediction. The gate mechanism is applied.

4.4.1 Evaluation Metrics

The competition used multi-class log-loss as evaluation metrics.

³DGL official website: <https://www.dgl.ai/pages/about.html>.

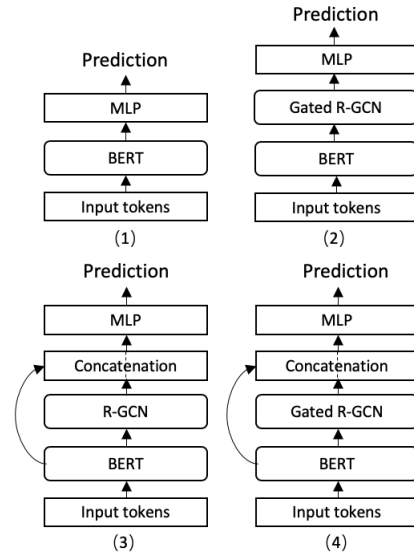


Figure 3: Stage one experiments

$$\text{logloss} = -\frac{1}{N} \sum_{i=1}^N \sum_{j=1}^M y_{ij} \log(p_{ij}),$$

where N is the number of samples in the test set, M is 3, \log is the natural logarithm.

4.4.2 Results

Table 1 presents the results of four different settings. it demonstrates that R-GCN structure does learn better embeddings and improve the performance. Setting three and setting four show the effectiveness of the Gate Mechanism.

BERT	R-GCN	Concatenation	Gate	Test Log-loss
Yes	No	No	No	0.5301
Yes	Yes	No	Yes	0.5142
Yes	Yes	Yes	No	0.5045
Yes	Yes	Yes	Yes	0.4936

Table 1: Stage one results

By comparing setting two and setting four, we can see that because graph convolution of the R-GCN model brings the potential problem of over-smoothing the information (Li et al., 2018), model without concatenation might lose some performance.

4.5 Full GAP Experiments and Results

We also tested our model on the full GAP dataset which contains 8,908 samples. 4908 samples were used as training data, and 4000 samples were used as test data. We used micro F1 score as our metric.

The GAP paper (Webster et al., 2018) introduced several baseline methods: (1) Off-the-shelf resolvers including a rule-based system of Lee et al. (2013) and three neural resolvers from Clark and Manning (2015), Wiseman et al. (2016), and Lee et al. (2017); (2) Baselines based on traditional cues for coreference; (3) Baselines based on structural cues: syntactic distance and Parallelism; (4) Baselines based on Wikipedia cues; (5) Transformer models (Vaswani et al., 2017).

Model	F1 Score
Lee et al. (2017)	64.0%
Parallelism	66.9%
Parallelism+URL	70.6%
BERT only	78.5%
Ours	80.3%

Table 2: GAP experiments results

Three best models (Lee et al. (2017), Parallelism, and Parallelism+URL) from above baselines were chosen for comparison. We first used pre-trained BERT embeddings and fully-connected layers for prediction (see Fig. 3 (1)). Not surprising, BERT embeddings outperformed all of the previous work.

We then tested our Gated R-GCN model. The model further improved the F1 score by using explicitly syntactic information and learning coreference-task-specific word representations. The final model largely increased the baseline F1 score from 70.6 % to 80.3 % and the BERT embeddings’ result from 78.5 % to 80.3 %.

4.6 Final Submission and Shared-Task Score

For the final submission for stage 2 of the shared task, we averaged our result with a BERT-score-layer (Zhang et al., 2018; Clark and Manning, 2016) result. In stage two, our work reaches log-loss of 0.394 on the private leaderboard showing that our model is quite effective and robust. This result is obtained without any data augmentation preprocessing.

5 Discussion and Conclusion

The Gender Bias for Natural Language Processing (GeBNLP) 2019 shared-task is a competition for building a coreference resolution system on GAP dataset. We participate in this shared-task by using a novel approach which combines Gated R-GCN with BERT. R-GCN is used for digesting syntactic

dependency graph and leveraging this syntactic information to help our semantic task. Experiments with four settings were conducted on the shared task’s stage one data. We also tested our model on the full GAP dataset where our model improved the best snippet-context baseline F1 score from 66.9 % to 80.3 % (by 20 %). The results showed that, under explicit syntactic supervision and without the need to fine tune BERT, our gated R-GCN model can incorporate syntactic structure prior with BERT embeddings to improve the performance on the coreference task.

References

- Kevin Clark and Christopher D Manning. 2015. Entity-centric coreference resolution with model stacking. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, volume 1, pages 1405–1415.
- Kevin Clark and Christopher D Manning. 2016. Deep reinforcement learning for mention-ranking coreference models. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2256–2262.
- Yann N Dauphin, Angela Fan, Michael Auli, and David Grangier. 2017. Language modeling with gated convolutional networks. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 933–941. JMLR. org.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- David Duvenaud, Dougal Maclaurin, Jorge Aguilera-Iparraguirre, Rafael Gómez-Bombarelli, Timothy Hirzel, Alán Aspuru-Guzik, and Ryan P Adams. 2015. Convolutional networks on graphs for learning molecular fingerprints. In *Proceedings of the 28th International Conference on Neural Information Processing Systems-Volume 2*, pages 2224–2232. MIT Press.
- Eraldo Rezende Fernandes, Cícero Nogueira dos Santos, and Ruy Luiz Milidiú. 2014. Latent trees for coreference resolution. *Computational Linguistics*, 40(4):801–835.
- Aria Haghighi and Dan Klein. 2009. Simple coreference resolution with rich syntactic and semantic features. In *Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing: Volume 3-Volume 3*, pages 1152–1161. Association for Computational Linguistics.

- Aria Haghighi and Dan Klein. 2010. Coreference resolution in a modular, entity-centered model. In *Human Language Technologies: The 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics*, pages 385–393. Association for Computational Linguistics.
- Jerry R Hobbs. 1978. Resolving pronoun references. *Lingua*, 44(4):311–338.
- Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Thomas N Kipf and Max Welling. 2016. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*.
- Shalom Lappin and Herbert J Leass. 1994. An algorithm for pronominal anaphora resolution. *Computational linguistics*, 20(4):535–561.
- Heeyoung Lee, Angel Chang, Yves Peirsman, Nathanael Chambers, Mihai Surdeanu, and Dan Jurafsky. 2013. Deterministic coreference resolution based on entity-centric, precision-ranked rules. *Computational Linguistics*, 39(4):885–916.
- Kenton Lee, Luheng He, Mike Lewis, and Luke Zettlemoyer. 2017. End-to-end neural coreference resolution. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 188–197.
- Qimai Li, Zhichao Han, and Xiao-Ming Wu. 2018. Deeper insights into graph convolutional networks for semi-supervised learning. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- Yujia Li, Daniel Tarlow, Marc Brockschmidt, and Richard Zemel. 2015. Gated graph sequence neural networks. *arXiv preprint arXiv:1511.05493*.
- Diego Marcheggiani and Ivan Titov. 2017. Encoding sentences with graph convolutional networks for semantic role labeling. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1506–1515.
- Joseph F McCarthy and Wendy G Lehnert. 1995. Using decision trees for coreference resolution. *arXiv preprint cmp-lg/9505043*.
- Sameer Pradhan, Alessandro Moschitti, Nianwen Xue, Olga Uryupina, and Yuchen Zhang. 2012. Conll-2012 shared task: Modeling multilingual unrestricted coreference in ontonotes. In *Joint Conference on EMNLP and CoNLL-Shared Task*, pages 1–40. Association for Computational Linguistics.
- Michael Schlichtkrull, Thomas N Kipf, Peter Bloem, Rianne Van Den Berg, Ivan Titov, and Max Welling. 2018. Modeling relational data with graph convolutional networks. In *European Semantic Web Conference*, pages 593–607. Springer.
- Dario Stojanovski and Alexander Fraser. 2018. Coreference and coherence in neural machine translation: A study using oracle experiments. In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 49–60.
- Rhea Sukthanker, Soujanya Poria, Erik Cambria, and Ramkumar Thirunavukarasu. 2018. Anaphora and coreference resolution: A review. *arXiv preprint arXiv:1805.11824*.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008.
- Kellie Webster, Marta Recasens, Vera Axelrod, and Jason Baldridge. 2018. Mind the gap: A balanced corpus of gendered ambiguous. In *Transactions of the ACL*, page to appear.
- Sam Wiseman, Alexander M Rush, and Stuart M Shieber. 2016. Learning global features for coreference resolution. *arXiv preprint arXiv:1604.03035*.
- Rui Zhang, Cicero Nogueira dos Santos, Michihiro Yasunaga, Bing Xiang, and Dragomir Radev. 2018. Neural coreference resolution with deep biaffine attention by joint mention detection and mention clustering. *arXiv preprint arXiv:1805.04893*.