

Towards Minimal Supervision BERT-Based Grammar Error Correction (Student Abstract)

Yiyuan Li, Antonios Anastasopoulos, Alan W Black

Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh, Pennsylvania 15213
yiyuanli@andrew.cmu.edu {aanastas, awb}@cs.cmu.edu

Abstract

Current grammatical error correction (GEC) models typically consider the task as sequence generation, which requires large amounts of annotated data and limit the applications in data-limited settings. We try to incorporate contextual information from pre-trained language model to leverage annotation and benefit multilingual scenarios. Results show strong potential of Bidirectional Encoder Representations from Transformers (BERT) in grammatical error correction task.

Introduction

The goal of grammatical error correction is to detect and correct all errors in the sentence and return the corrected sentences. Current grammatical error correction approaches require a large amount of training data to achieve reasonable results, which unfortunately cannot be extended to languages with limited data. Recently, unsupervised models pre-trained on large corpora have boosted performance in many natural language processing tasks, which indicates a potential of leveraging such models for GEC in any language. In this work, we try to utilize (multilingual) BERT (Devlin et al. 2019) in order to perform grammatical error correction with minimum supervision.

Proposed Method

Our approach divides the GEC task into two stages: error identification followed by correction. In the first stage, we try to detect the span in the original text that the edit will apply. We formulate this as a sequence labelling task where tokens are labelled in one of the following labels {*remain*, *substitution*, *insert*, *delete*}. For the second stage (correction) we employ a pretrained-model like BERT. The labels from the error identification stage guide the masking of the inputs (where we mask any tokens marked with *substitution* or insert mask tokens for *insert* labels), and we obtain candidate outputs for every masked token. A shortcoming of our current approach is that it only produces as many corrections as masked input tokens; however, most grammar errors in fact tend to be edits of length 1 or 2, which are captured

by our identification labels. The second stage is addressed in future work.

Results

We report preliminary results on the test part of the English FCE dataset (Yannakoudakis, Briscoe, and Medlock 2011). The edits are labelled and scored through ERRANT (Bryant, Felice, and Briscoe 2017).

For our preliminary experiments we focus on a simplified single-edit setting, where we attempt to correct sentences with a single error (assuming oracle error identification annotations). The goal is to assess the capabilities of pre-trained BERT-like models assuming perfect performance for all other components (e.g. error identification).

We expand the original dataset to fit our single-error scenario with the following two schemes. (1) **each edit**: all corrections except one are applied, creating a single-error sentence; (2) **last edit**: all corrections except for the last one are applied. After the split, we obtain 3,585 and 1,024 corrections respectively. We also employ different strategies for deciding the number of masked tokens: (1) based on span length of the original edit, (2) based on length of the final correction (given from an oracle), and (3) using a single mask and measure whether any token of the correction is predicted. Note that subword predictions like {*ad*, *##e*, *##quote*} to adequate are merged in sentence-level, but remains in mask-level evaluation. Our preliminary results under the various settings are outlined in Table 1. We find that different masking strategies have comparable performance, with slightly higher accuracy when using a single mask. Interestingly, BERT seems capable to actually produce corrections with quite high precision of more than 70%.

We also study whether the correct output is among the top-*k* candidates suggested by BERT. We compare the $F_{0.5}$ metrics based on each mask between the most probable prediction and the top *k* candidates, where *k* is set to 5. From the result in Table 2, we observe that an appropriate reranking model could further boost the performance by selecting the appropriate correction.

Masking Strategy	each edit			last edit		
	P@1	R@1	F _{0.5} @1	P@1	R@1	F _{0.5} @1
# origin	0.632	0.853	0.667	0.592	0.824	0.627
# target	0.66	0.887	0.696	0.614	0.855	0.651
single	0.763	0.931	0.790	0.767	0.920	0.794

Table 1: Sentence-level evaluation with different masking strategies for single edit pairs. Subword predictions are merged in sentence generation.

Masking Strategy	each edit		last edit	
	Acc@1	Acc@5	Acc@1	Acc@5
# origin	0.292	0.455	0.229	0.390
# target	0.313	0.484	0.247	0.405
single	0.365	0.554	0.312	0.501

Table 2: A reranking mechanism could lead to better results, as performance@5 is higher than performance@1.

Related and Future work

A retrieve-edit model is proposed for text generation (Guu et al. 2017). However, the edition is one-time and sentences with multiple grammatical errors could further reduce the similarity between the correct form and the oracle sentence. An iterative decoding approach (Ge, Wei, and Zhou 2018) or the neural language model (Choe et al. 2019) as the scoring function are employed for GEC. To the best of our knowledge, there is no prior work in applying pre-trained contextual model in grammatical error correction. In the future work, we will additionally model *error fertility*, allowing us to exactly predict the number of necessary [MASK] tokens. Last, we will employ a re-ranking mechanism which scores the candidate outputs from BERT, taking into account larger context and specific grammar properties.

Better span detection Although BERT could predict all the missing token in the sentence in a reasonable way, prediction of the correct words could easily fall into redundant editing. Our experiment shows that simply rephrasing the whole sentence using BERT would lead to too diverse an output. Instead, a prior error span detection could be necessary for efficient GEC, and it is part of our future work.

Partial masking and fluency measures Multi-masking or masking an informative part in the sentence will lead to loss of original information, and it will allow unwanted freedom in the predictions; see Table 3 for examples. Put in plain terms, multi-masking allows BERT to get too creative. Instead, we will investigate partial masking strategies (Zhou et al. 2019) which could alleviate this problem. Fluency is an important measure when employing an iterative approach (Napoles, Sakaguchi, and Tetreault 2016). We plan to explore fluency measures as part of our reranking mechanisms.

Example #1: Redundant Edits	
Source	Of course there 's also a number 8 bus in front of the hotel , which is also suitable , but it leaves only every half an hour
Mask.	Of course there 's also a number 8 bus [MASK] in front of the hotel , which is also suitable , but it leaves only every half an hour
Target	Of course there 's also a number 8 bus , in front of the hotel , which is also suitable , but it leaves only every half an hour
Ours	Of course there 's also a number 8 bus <u>stop</u> in front of the hotel , which is also suitable , but it leaves only every half an hour
Example #2 : Synonyms	
Source	The aim of this report is to <u>recomend</u> you to visit the Fuerte de San Diego Museum
Mask.	The aim of this report is to [MASK] you to visit the Fuerte de San Diego Museum
Target	The aim of this report is to <u>recommend</u> you to visit the Fuerte de San Diego Museum
Ours	The aim of this report is to <u>allow</u> you to visit the Fuerte de San Diego Museum
Example #3 : Hallucination	
Source	Of course there 's also a <u>bus number 8</u> , in front of the hotel , which is also suitable , but it leaves only every half an hour
Mask.	Of course there 's also a [MASK] [MASK] [MASK], in front of the hotel , which is also suitable , but it leaves only every half an hour
Target	Of course there 's also a <u>number 8 bus</u> , in front of the hotel , which is also suitable , but it leaves only every half an hour
Ours	Of course there 's also a small <u>parking station</u> , in front of the hotel , which is also suitable , but it leaves only every half an hour

Table 3: Common BERT prediction errors (we highlight the original error and the prediction).

References

- Bryant, C.; Felice, M.; and Briscoe, T. 2017. Automatic annotation and evaluation of error types for grammatical error correction. In *Proc. ACL*.
- Choe, Y. J.; Ham, J.; Park, K.; and Yoon, Y. 2019. A neural grammatical error correction system built on better pre-training and sequential transfer learning. In *Proc. BEA@ACL*.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proc. NAACL-HLT*.
- Ge, T.; Wei, F.; and Zhou, M. 2018. Fluency boost learning and inference for neural grammatical error correction. In *Proc. ACL*.
- Guu, K.; Hashimoto, T. B.; Oren, Y.; and Liang, P. 2017. Generating sentences by editing prototypes. *TACL* 6:437–450.
- Napoles, C.; Sakaguchi, K.; and Tetreault, J. 2016. There's no comparison: Reference-less evaluation metrics in grammatical error correction. In *Proc. EMNLP*, 2109–2115.
- Yannakoudakis, H.; Briscoe, T.; and Medlock, B. 2011. A new dataset and method for automatically grading ESOL texts. In *Proc. ACL*.
- Zhou, W.; Ge, T.; Xu, K.; Wei, F.; and Zhou, M. 2019. BERT-based lexical substitution. In *Proc. ACL*.