

Improving the Efficiency of Grammatical Error Correction with Erroneous Span Detection and Correction

Mengyun Chen^{1*} Tao Ge^{2†} Xingxing Zhang² Furu Wei² Ming Zhou²

¹ East China Normal University; ² Microsoft Research Asia

dreamcloud.chen@gmail.com;

{tage, xizhang, fuwei, mingzhou}@microsoft.com

Abstract

We propose a novel language-independent approach to improve the efficiency for Grammatical Error Correction (GEC) by dividing the task into two subtasks: Erroneous Span Detection (ESD) and Erroneous Span Correction (ESC). ESD identifies grammatically incorrect text spans with an efficient sequence tagging model. Then, ESC leverages a seq2seq model to take the sentence with annotated erroneous spans as input and only outputs the corrected text for these spans. Experiments show our approach performs comparably to conventional seq2seq approaches in both English and Chinese GEC benchmarks with less than 50% time cost for inference.

1 Introduction

Due to a growing number of error-corrected parallel sentences available in recent years, sequence-to-sequence (seq2seq) models with the encoder-decoder architecture (Bahdanau et al., 2014; Sutskever et al., 2014; Luong et al., 2015) have become a popular solution to GEC, which take the source (original) sentence as input and output the target (corrected) sentence. Although autoregressive seq2seq models facilitate correction for various grammatical errors and perform well, they are not efficient enough for GEC. As previous work (Zhao et al., 2019) points out, seq2seq models take most decoding steps to copy grammatically correct text spans from the source to the target during inference, which is the main efficiency bottleneck. If the time for the copying operations can be saved, the efficiency should be much improved.

With this motivation, we propose a simple yet novel language-independent approach to improve the efficiency of GEC by dividing the task into

two subtasks: Erroneous Span Detection (ESD) and Erroneous Span Correction (ESC), shown in Figure 1. In ESD, we use an efficient sequence tagging model to identify the text spans that are grammatically incorrect in the source sentence, as Figure 1(a) shows. Then, we feed the sentence with erroneous span annotations to a seq2seq model for ESC. In contrast to conventional seq2seq approaches correcting the complete sentence, ESC only corrects the erroneous spans (see Figure 1(b)), which largely decreases the number of steps for decoding. Experiments in both English and Chinese GEC benchmarks demonstrate our approach performs comparably to the state-of-the-art transformer based seq2seq model with less than 50% time cost for inference. Furthermore, our approach offers more flexibility to control correction, allowing us to adapt the precision-recall trade-off to various application scenarios.

2 Related Work

Recently, many approaches have been proposed to improve GEC performance. However, except those adding synthetic erroneous data (Xie et al., 2018; Ge et al., 2018a; Grundkiewicz et al., 2019; Kiyono et al., 2019; Zhou et al., 2019) and Wikipedia revision logs (Lichtarge et al., 2019) for training, most methods cause an increase in latency. For example, language model and right-to-left (R2L) rescoring (Grundkiewicz et al., 2019; Kiyono et al., 2019) not only take time to rescore but also slow down the correction model with a larger beam size during inference; multi-round (iterative) decoding (Ge et al., 2018a,b; Lichtarge et al., 2019) needs to repeatedly run the model; BERT-fuse (Kaneko et al., 2020) adds extra computation for model fusion.

In contrast to extensive studies on GEC performance, little work focuses on improving the efficiency of GEC models until the last years.

* This work was done during the author’s internship at Microsoft Research Asia.

† Co-first authors with equal contributions.

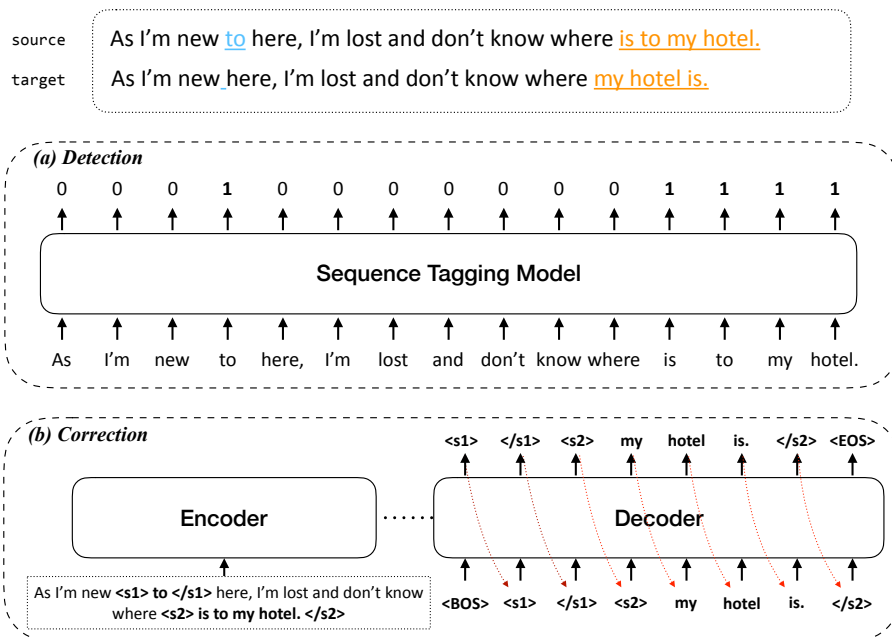


Figure 1: An overview of erroneous span detection (ESD) and erroneous span correction (ESC). The detection model is a sequence tagging model, while the correction model is a seq2seq model but only outputs the corrected texts for annotated spans (i.e., $\langle s1 \rangle$ to $\langle /s1 \rangle$ and $\langle s2 \rangle$ is to my hotel. $\langle /s2 \rangle$). For example, the text span “**is to my hotel.**” is identified as incorrect and then edited into “**my hotel is.**”.

One branch of the work is language-dependent approaches, like PIE (Awasthi et al., 2019) and GECToR (Omelianchuk et al., 2020). They predict a sequence of token-level edit operations including a number of manually designed language-specific operations like changing verb forms (e.g., $VBZ \rightarrow VBD$) and prepositions (e.g., $in \rightarrow on$). However, they are difficult to be adapted to other languages. The other branch is language-independent models like LaserTagger (Malmi et al., 2019). They learn a vocabulary of edit operations from training data and thus can work for any language. However, their performance is inferior to their seq2seq counterpart. Our approach combines the advantages of both branches, which is language-independent and performs comparably to the state-of-the-art seq2seq approach with efficient inference.

3 Erroneous Span Detection

To identify incorrect spans, we use a binary sequence tagging model in which tag **0** means the token is in a correct span; while tag **1** means the token is in a grammatically incorrect span that needs to be edited, as shown in Figure 1(a). In order to train the tagging model, we align¹ tokens across the

¹Alignment can be solved by dynamic programming like Levenshtein distance. We here use ERRANT (<https://github.com/chrisjbryant/errant>) for alignment.

source and target sentence in training data. With token alignment, we can identify the text spans that are edited and thus can annotate the edited text spans in the original sentences as erroneous spans.

4 Erroneous Span Correction

With ESD, we can identify grammatically incorrect text spans in a sentence. If a sentence is identified as error-free, we take no further action; otherwise, we annotate the incorrect spans and use the ESC model to correct them, shown in Figure 1(b).

To avoid ESC being misled by span detection errors from ESD during inference, we randomly select text spans in the similar way to SpanBERT (Joshi et al., 2019) instead of only using gold erroneous spans in training data, to train the ESC model. In this way, the ESC model will see a large variety of span annotations and learn how to correct during training, and thus its robustness is improved: even if the detected spans during inference are not exactly accurate, the ESC model will not easily fail. With token alignment across the source and target sentence in GEC training data, we can generate training instances with span annotations and corrections like the example in Figure 1(b) for ESC.

Model	Pretrained	Faster/Slower	CoNLL-14 (M^2)			BEA-19 (ERRANT)		
			P	R	$F_{0.5}$	P	R	$F_{0.5}$
<i>Seq2seq</i>	No	-	64.9	26.6	50.4	57.3	41.5	53.2
<i>Levenshtein Transformer*</i> (Gu et al., 2019)	No	Faster	39.9	24.4	35.4	32.2	39.2	33.4
<i>Levenshtein Transformer*</i> (distillation)	No	Faster	53.1	23.6	42.5	45.5	37.0	43.5
<i>LaserTagger*</i> (Malmi et al., 2019)	No	Faster	50.9	26.9	43.2	53.4	38.5	49.6
Our approach	No	Faster	66.0	24.7	49.5	62.7	38.6	55.7
<i>Seq2seq</i>	Yes	-	69.4	42.5	61.5	66.7	61.3	65.5
<i>PRETLarge</i> (Kiyono et al., 2019)	Yes	-	67.9	44.1	61.3	65.5	59.4	64.2
<i>PIE</i> (Awasthi et al., 2019)	Yes	Faster	66.1	43.0	59.7	<u>58.0</u>	<u>53.1</u>	<u>56.9</u>
Our approach	Yes	Faster	72.6	37.2	61.0	70.4	55.9	66.9
<i>BERT-fuse GED</i> (Kaneko et al., 2020)	Yes	Slower	69.2	45.6	62.6	67.1	60.1	65.6
<i>BERT-fuse GED+R2L</i> (Kaneko et al., 2020)	Yes	Slower	72.6	46.4	65.2	72.3	61.4	69.8
<i>PRETLarge+SSE+R2L</i> (Kiyono et al., 2019)	Yes	Slower	72.4	46.1	65.0	72.1	61.8	69.8
<i>UEDIN-MS</i> (Grundkiewicz et al., 2019)	Yes	Slower	-	-	64.2	72.3	60.1	69.5

Table 1: Performance in English GEC benchmarks (i.e., CoNLL-14 and BEA-19 test). *Seq2seq* is our implemented seq2seq model based on **Transformer (big)** architecture, which is also the baseline for speed comparison (i.e., **Faster/Slower** in the table). The column **Pretrained** indicates whether the model is pretrained with synthetic or additional (e.g., Wikipedia revision logs) error-corrected data. * indicates the models are implemented by us with the released codes of the original papers, trained and evaluated on the BEA-19 setting. The underlines indicate the scores are evaluated by us for the released model on the BEA-19 test data.

Model	CoNLL-14 (1,312)							NLPCC-18 (2,000)						
	Time (in second)				Performance			Time (in second)				Performance		
	1	8	16	32	P	R	$F_{0.5}$	1	8	16	32	P	R	$F_{0.5}$
<i>Seq2seq</i>	363	85	51	33	64.9	26.6	50.4	690	166	101	63	36.9	14.4	28.1
<i>Levenshtein Transformer</i>	125	24	19	14	53.1	23.6	42.5	224	64	41	31	24.9	15.0	22.0
Our approach	137	34	21	16	66.0	24.7	49.5	253	60	39	29	37.3	14.5	28.4
<i>Seq2seq (tensor2tensor)</i>	680	138	97	85	58.7	30.5	49.5	1292	227	141	92	41.0	10.8	26.3
<i>PIE</i>	66	52	51	48	-	-	-	-	-	-	-	-	-	-
<i>LaserTagger</i>	23	12	9	8	50.9	26.9	43.2	34	16	14	13	25.6	10.5	19.9

Table 2: Performance and total inference time of models without pretraining under various batch sizes (1/8/16/32) using 1 Nvidia V100 GPU with CUDA 10.2 in the English (CoNLL-14: 1,312 sentences) and Chinese (NLPCC-18: 2,000 sentences) GEC test sets. The top group of models is implemented with Pytorch, while the bottom group is implemented with Tensorflow, thus their inference time cannot be compared. The performance of PIE in CoNLL-14 is not reported because it is pretrained with synthetic data and thus unfair to be compared here. Also, PIE has no result in NLPCC-18 because it is specific for English and difficult to be generalized to other languages.

5 Experiments

5.1 Experimental Setting

Following recent work in English GEC, we conduct experiments in the same setting with the restricted track of the BEA-2019 GEC shared task (Bryant et al., 2019), using FCE (Yannakoudakis et al., 2011), Lang-8 Corpus of Learner English (Mizumoto et al., 2011), NUCLE (Dahlmeier et al., 2013) and W&I+LOCNESS (Granger, 1998; Bryant et al., 2019) as training data. We use CoNLL-2013 test set as the dev set to choose the best-performing models, and evaluate on the well-known GEC benchmark datasets: CoNLL-2014 (Ng et al., 2014) and BEA-2019 test set with the official evaluation scripts (m2scorer² for CoNLL-14, ERRANT for BEA-19). As previous work (Grundkiewicz et al., 2019) trained with synthetic data, we synthesize

260M sentence pairs in the same way to try pre-training ESD and ESC. Also, we verify in Chinese GEC whether our approach can be adapted to other languages. We follow the setting of NLPCC-2018 Chinese GEC shared task (Zhao et al., 2018), using its official training³ and evaluation datasets.

We fine-tune Roberta (Liu et al., 2019) base and Chinese Bert (Devlin et al., 2018) base model for English and Chinese ESD respectively. For ESC, we train a Transformer (big) model (Vaswani et al., 2017), using an encoder-decoder shared vocabulary of 32K Byte Pair Encoding (Sennrich et al., 2015) tokens for English and 8.4K Chinese character for Chinese. During inference, ESC decodes with a beam size of 5. We include more details of models, training and inference in the supplementary notes.

²<https://github.com/nusnlp/m2scorer>

³We sample 5,000 training instances as the dev set.

Model	#Sent for ESC/seq2seq	Batch size			P	R	$F_{0.5}$
		1	8	16			
ESD (base) + ESC	872	137 (23+114)	34 (5+29)	21 (3+18)	66.0	24.7	49.5
ESD (base) + seq2seq	872	284 (23+261)	66 (5+61)	41 (3+38)	67.0	25.1	50.2
ESD (large) + ESC	935	167 (43+124)	40 (8+32)	25 (6+19)	67.2	26.4	51.3
ESD (large) + seq2seq	935	318 (43+275)	77 (8+69)	47 (6+41)	66.6	25.5	50.3
seq2seq	1,312	363	85	51	64.9	26.6	50.4

Table 3: In-depth time cost (in second) analysis in CoNLL-14 which contains 1,312 test sentences. (base) and (large) indicate that the ESD models are fine-tuned from the Roberta base and large models respectively. ESD + seq2seq is implemented as follows: ESD first identifies the sentences that have grammatical errors, then the seq2seq baseline model only corrects these sentences. The column #Sent for ESC/seq2seq shows the actual number of sentences ESC/seq2seq processed. For the time cost in the brackets such as (23+114), the first term (e.g., 23) is the time cost by the ESD model while the last term is the cost (e.g., 114) by the other parts.




Figure 2: Attention Heatmap of ESC.

5.2 Experimental Results

Table 1 shows the performance of our approach and recent models in English GEC. Although the non auto-regressive models like Levenshtein Transformer and LaserTagger are faster than the seq2seq baseline, their performance is not desirable. Among the models without pretraining (**top group**), our approach is the only one that performs comparably to the *Seq2seq* baseline with faster inference. When we add the synthetic data to pretrain the ESD and ESC models, our approach’s results are much improved, yielding state-of-the-art results among the models with good inference speed (**middle group**). Though our approach underperforms the best systems (**bottom group**) which improve results through various methods (e.g., model fusion, ensemble decoding and rescoring) that seriously hurt efficiency, it is much more efficient and thus applicable in real world applications.

Table 2 compares the inference time of various approaches. Compared to the *Seq2seq* implementation in Pytorch-fairseq, our approach saves over 50% time cost. It is notable that among the implementations in Table 2, LaserTagger is the most efficient though its results are not good enough.

threshold	0.2	0.3	0.4	0.5	0.6	0.7
P	62.4	63.8	64.8	66.0	66.2	67.0
R	27.4	26.8	25.6	24.7	23.4	21.7
$F_{0.5}$	49.7	50.0	49.6	49.5	48.4	47.2

Table 4: As the probability threshold of ESD increases, precision increases while recall drops in CoNLL-14.

Also, our approach consistently achieves comparable performance with the *Seq2seq* baseline in Chinese GEC, demonstrating that our approach can be easily adapted to other languages.

We further analyze the corresponding time cost of ESD and ESC. Table 3 shows that ESD is much faster than the auto-regressive ESC model. It not only efficiently filters out error-free sentences to save effort for the following process, but also pinpoints incorrect spans, allowing ESC to only focus on correcting the spans and reduce the decoding steps, as shown in Figure 2. For the 872 sentences ESD (base) identifies as incorrect, the total number of decoding time steps (in the best beam) of ESC is 7,647, accounting for the efficiency improvement over the seq2seq model whose corresponding decoding steps are 21,065. Furthermore, if we use a larger ESD model (Roberta large), we observe better results with a still marked reduction in time cost compared to the baseline. More detailed qualitative studies and analyses of ESD are presented in the supplementary notes due to the space limitation.

Besides its efficiency advantage, our approach offers more flexibility to control correction behavior during inference, making it adaptive to various real-world application scenarios. As shown in Table 4, if the model is intended for high precision, we can increase the probability threshold for ESD so that it identifies the incorrect spans only when it is very confident; on the other hand, if we want the model to be aggressive for higher recall, we can simply decrease the threshold.

6 Conclusion and Future Work

We propose a novel language-independent approach to improve the efficiency of GEC. Our approach performs comparably to the state-of-the-art seq2seq model with a considerable reduction in inference time, and can be easily adapted to other languages and offer more flexibility to control correction behavior (e.g., trading precision for recall).

Through our experiments in GEC, we verify the feasibility of span-specific decoding, which has been explored for text infilling (Raffel et al., 2019) and text rewriting. It is inspiring and promising to be generalized to more rewriting tasks, which will be studied as our future work.

Acknowledgments

We thank Xin Sun and Xiangxin Zhou for their help with experiments. We thank the anonymous reviewers for their valuable comments. The corresponding author of this paper is Tao Ge.

References

- Abhijeet Awasthi, Sunita Sarawagi, Rasna Goyal, Sabyasachi Ghosh, and Vihari Piratla. 2019. Parallel iterative edit models for local sequence transduction. *arXiv preprint arXiv:1910.02893*.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*.
- Christopher Bryant, Mariano Felice, Øistein E Andersen, and Ted Briscoe. 2019. The bea-2019 shared task on grammatical error correction. In *Proceedings of the Fourteenth Workshop on Innovative Use of NLP for Building Educational Applications*, pages 52–75.
- Daniel Dahlmeier, Hwee Tou Ng, and Siew Mei Wu. 2013. Building a large annotated corpus of learner english: The nus corpus of learner english. In *Proceedings of the eighth workshop on innovative use of NLP for building educational applications*, pages 22–31.
- 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*.
- Sergey Edunov, Myle Ott, Michael Auli, and David Grangier. 2018. Understanding back-translation at scale. *arXiv preprint arXiv:1808.09381*.
- Tao Ge, Furu Wei, and Ming Zhou. 2018a. Fluency boost learning and inference for neural grammatical error correction. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1055–1065.
- Tao Ge, Furu Wei, and Ming Zhou. 2018b. Reaching human-level performance in automatic grammatical error correction: An empirical study. *arXiv preprint arXiv:1807.01270*.
- Sylviane Granger. 1998. *The computer learner corpus: a versatile new source of data for SLA research*. na.
- Roman Grundkiewicz, Marcin Junczys-Dowmunt, and Kenneth Heafield. 2019. Neural grammatical error correction systems with unsupervised pre-training on synthetic data. In *Proceedings of the Fourteenth Workshop on Innovative Use of NLP for Building Educational Applications*, pages 252–263.
- Jiatao Gu, Changhan Wang, and Junbo Zhao. 2019. Levenshtein transformer. In *Advances in Neural Information Processing Systems*, pages 11179–11189.
- Mandar Joshi, Danqi Chen, Yinhan Liu, Daniel S. Weld, Luke Zettlemoyer, and Omer Levy. 2019. SpanBERT: Improving pre-training by representing and predicting spans. *arXiv preprint arXiv:1907.10529*.
- Masahiro Kaneko, Masato Mita, Shun Kiyono, Jun Suzuki, and Kentaro Inui. 2020. Encoder-decoder models can benefit from pre-trained masked language models in grammatical error correction. *arXiv preprint arXiv:2005.00987*.
- Masahiro Kaneko, Yuya Sakaizawa, and Mamoru Komachi. 2017. Grammatical error detection using error- and grammaticality-specific word embeddings. In *Proceedings of the Eighth International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*.
- Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Shun Kiyono, Jun Suzuki, Masato Mita, Tomoya Mizumoto, and Kentaro Inui. 2019. An empirical study of incorporating pseudo data into grammatical error correction. *arXiv preprint arXiv:1909.00502*.
- Jared Lichtarge, Chris Alberti, Shankar Kumar, Noam Shazeer, Niki Parmar, and Simon Tong. 2019. Corpora generation for grammatical error correction. *arXiv preprint arXiv:1904.05780*.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [Roberta: A robustly optimized BERT pretraining approach](#). *CoRR*, abs/1907.11692.
- Minh-Thang Luong, Hieu Pham, and Christopher D Manning. 2015. Effective approaches to attention-based neural machine translation. *arXiv preprint arXiv:1508.04025*.

- Eric Malmi, Sebastian Krause, Sascha Rothe, Daniil Mirylenka, and Aliaksei Severyn. 2019. Encode, tag, realize: High-precision text editing. *arXiv preprint arXiv:1909.01187*.
- Tomoya Mizumoto, Mamoru Komachi, Masaaki Nagata, and Yuji Matsumoto. 2011. Mining revision log of language learning sns for automated japanese error correction of second language learners. In *Proceedings of 5th International Joint Conference on Natural Language Processing*, pages 147–155.
- Courtney Napoles, Keisuke Sakaguchi, Matt Post, and Joel Tetreault. 2015. Ground truth for grammatical error correction metrics. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 588–593.
- Courtney Napoles, Keisuke Sakaguchi, and Joel Tetreault. 2017. Jfleg: A fluency corpus and benchmark for grammatical error correction. *arXiv preprint arXiv:1702.04066*.
- Hwee Tou Ng, Siew Mei Wu, Ted Briscoe, Christian Hadiwinoto, Raymond Hendy Susanto, and Christopher Bryant. 2014. The conll-2014 shared task on grammatical error correction. In *Proceedings of the Eighteenth Conference on Computational Natural Language Learning: Shared Task*, pages 1–14.
- Kostiantyn Omelianchuk, Vitaliy Atrasevych, Artem Chernodub, and Oleksandr Skurzhanyskiy. 2020. Gector–grammatical error correction: Tag, not rewrite. *arXiv preprint arXiv:2005.12592*.
- Robert Parker, David Graff, Junbo Kong, Ke Chen, and Kazuaki Maeda. 2011. English gigaword fifth edition, linguistic data consortium. *Google Scholar*.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2019. Exploring the limits of transfer learning with a unified text-to-text transformer. *arXiv preprint arXiv:1910.10683*.
- Marek Rei. 2017. Semi-supervised multitask learning for sequence labeling. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*.
- Marek Rei and Helen Yannakoudakis. 2017. Auxiliary objectives for neural error detection models. In *Proceedings of the 12th Workshop on Innovative Use of NLP for Building Educational Applications*.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2015. Neural machine translation of rare words with subword units. *arXiv preprint arXiv:1508.07909*.
- Ilya Sutskever, Oriol Vinyals, and Quoc V Le. 2014. Sequence to sequence learning with neural networks. In *Advances in neural information processing systems*, pages 3104–3112.
- Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. 2016. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2818–2826.
- 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.
- Ziang Xie, Guillaume Genthial, Stanley Xie, Andrew Y Ng, and Dan Jurafsky. 2018. Noising and denoising natural language: Diverse backtranslation for grammar correction. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 619–628.
- Helen Yannakoudakis, Ted Briscoe, and Ben Medlock. 2011. A new dataset and method for automatically grading esol texts. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies-Volume 1*, pages 180–189. Association for Computational Linguistics.
- Yi Zhang, Tao Ge, Furu Wei, Ming Zhou, and Xu Sun. 2019. Sequence-to-sequence pre-training with data augmentation for sentence rewriting. *arXiv preprint arXiv:1909.06002*.
- Wei Zhao, Liang Wang, Kewei Shen, Ruoyu Jia, and Jingming Liu. 2019. Improving grammatical error correction via pre-training a copy-augmented architecture with unlabeled data. *arXiv preprint arXiv:1903.00138*.
- Yuanyuan Zhao, Nan Jiang, Weiwei Sun, and Xiaojun Wan. 2018. Overview of the nlpcc 2018 shared task: grammatical error correction. In *CCF International Conference on Natural Language Processing and Chinese Computing*, pages 439–445. Springer.
- Wangchunshu Zhou, Tao Ge, Chang Mu, Ke Xu, Furu Wei, and Ming Zhou. 2019. Improving grammatical error correction with machine translation pairs. *arXiv preprint arXiv:1911.02825*.

A Experiment Details

Table 5 describes the details of datasets used for English GEC. Except the sythetic data, all the data can be found at the website⁴ of the BEA-19 shared task. The sythetic data is generated from English Wikipedia⁵, English Gigaword (Parker et al., 2011) and Newscrawl⁶ as the previous work (Ge et al.,

⁴<https://www.cl.cam.ac.uk/research/nlp/bea2019st/>

⁵<https://en.wikipedia.org/>

⁶<http://data.statmt.org/news-crawl/en/>

Corpus	Sent.	Tok.	Usage
Synthetic data	256M	4.88B	Pretraining
FCE	32.0K	0.5M	Fine-tuning
Lang-8	1.04M	13.0M	Fine-tuning
NUCLE	57.1K	1.2M	Fine-tuning
W&I+LOCNESS	34.3K	0.6M	Fine-tuning
CoNLL-13	1,381	29.2K	Development
BEA-19	4,477	89.3K	Test
CoNLL-14	1,312	30.1K	Test

Table 5: Statistics of the datasets used for pretraining, fine-tuning and evaluation.

Configurations	Values
Pretraining	
Model Architecture	Roberta (base) (Liu et al., 2019)
Number of parameters	125M
Number of epochs	5
Devices	8 Nvidia V100 GPU
Max tokens per GPU	12000
Optimizer	Adam ($\beta_1=0.9, \beta_2=0.98, \epsilon=1 \times 10^{-6}$) (Kingma and Ba, 2014)
Learning rate	3×10^{-5}
Learning rate scheduler	inverse sqrt
warmup	8000
weight decay	0.1
Dropout	0.3
Fine-tuning	
Number of epochs	50
Devices	4 Nvidia V100 GPU
Max tokens per GPU	8000
Learning rate	1×10^{-5}
warmup	4000
Dropout	0.2

Table 6: Hyper-parameters values of ESD during the pretraining and fine-tuning.

2018a; Zhang et al., 2019; Kiyono et al., 2019; Grundkiewicz et al., 2019) did, using back translation and sentence corruption. Specifically, we train a transformer (base) model (Vaswani et al., 2017) for back translation using the training data of the restricted track in the BEA-19 shared task. For sentence corruption, we follow Edunov et al. (2018) to randomly insert, delete, replace and swap adjacent tokens in a sentence.

Hyper-parameters for the ESD and ESC model for English GEC are listed in table 6 and table 7. The hyper-parameters for Chinese GEC are almost the same except that the ESD model is fine-tuned from Chinese Bert (Devlin et al., 2018) base model.

At last, we highlight that the Levenshtein Transformer baselines in this paper are implemented using the master branch of fairseq (previous versions may have different reproduced results).

Configurations	Values
Pretraining	
Model Architecture	Transformer (big) (Vaswani et al., 2017)
Number of parameters	209M
Number of epochs	5
Devices	8 Nvidia V100 GPU
Max tokens per GPU	12000
Optimizer	Adam ($\beta_1=0.9, \beta_2=0.98, \epsilon=1 \times 10^{-8}$) (Kingma and Ba, 2014)
Learning rate	5×10^{-4}
Learning rate scheduler	polynomial decay
Warmup	8000
weight decay	0.0
Loss Function	label smoothed cross entropy (label-smoothing=0.1) (Szegedy et al., 2016)
Dropout	0.3
Fine-tuning	
Number of epochs	30
Devices	4 Nvidia V100 GPU
Max tokens per GPU	5120
Learning rate	3×10^{-4}
Warmup	4000
Beam search	5

Table 7: Hyper-parameters values of ESC during the pretraining and fine-tuning.

ESD Model	Annotation 1			Annotation 2		
	P	R	$F_{0.5}$	P	R	$F_{0.5}$
ESD (base)	52.4	35.3	47.8	63.4	31.3	52.6
ESD (large)	49.5	40.0	47.2	61.4	36.4	54.0
ESD (base+pretrained)	50.9	40.1	48.3	63.7	36.8	55.6

Table 8: The performance of ESD on the two official annotations for the CoNLL-14 shared task test dataset.

B ESD performance

We followed the previous work (Rei, 2017; Rei and Yannakoudakis, 2017; Kaneko et al., 2017) in Grammatical Error Detection (GED), using token-level *precision*, *recall* and $F_{0.5}$ to evaluate our ESD model. Table 8 shows the results in CoNLL-14.

C Examples

In table 9 and table 10, we show examples that are corrected by our approach to demonstrate the effectiveness of our approach in practice. According to the results in these tables, it is clear that our approach can yield satisfying corrections without hurting fluency, which is consistent with our evaluation results in the JFLEG (Napoles et al., 2017) test set with respect to GLEU⁷ (Napoles et al., 2015) – an automatic fluency metric for GEC.

⁷Our approach (without pretraining) achieves 53.0 GLEU, comparable to 52.7 by its seq2seq counterpart in JFLEG.

Source Sentence	Instead, we will post a seed and tag our friends to inform this kind of changments.
Annotation	Instead, we will post a seed and tag our friends to <s1> inform this kind of changments. </s1>
Correction	<s1> inform them of these kinds of changes. </s1>
Final Output	Instead, we will post a seed and tag our friends to inform them of these kinds of changes.
Source Sentence	Personally I feel that we still should take our responsibility to tell them the situation.
Annotation	<s1> Personally I </s1> feel that <s2> we still should </s2> take our responsibility to tell them the situation.
Correction	<s1> Personally , I </s1> <s2> we should still </s2>
Final Output	Personally, I feel that we should still take our responsibility to tell them the situation.
Source Sentence	The law 's spirit also include the fairness.
Annotation	The law 's spirit <s1> also include the fairness. </s1>
Correction	<s1> also includes fairness. </s1>
Final Output	The law 's spirit also includes fairness.
Source Sentence	Above all, life is more important than secret.
Annotation	Above all, life is more important <s1> than secret. </s1>
Correction	<s1> than a secret. </s1>
Final Output	Above all, life is more important than a secret.
Source Sentence	So, they have to also prepare mentally.
Annotation	So, <s1> they have to also prepare </s1> mentally.
Correction	<s1> they also have to prepare </s1>
Final Output	So, they also have to prepare mentally.
Source Sentence	To prevent the bigger problem to happen, it takes a lot of effort to take care of your body.
Annotation	To prevent the bigger <s1> problem to happen, it </s1> takes a lot of effort to take care of your body.
Correction	<s1> problem from happening, it </s1>
Final Output	To prevent the bigger problem from happening, it takes a lot of effort to take care of your body.

Table 9: Examples of our ESD & ESC approach in English for GEC. ESD first detects the grammatical incorrect text spans in the source sentence. Then the sentence with the erroneous span annotations (the **Annotation** row) are fed into the ESC model to generate the corresponding corrections (the **Correction** row) for the annotated spans. Finally, we replace the erroneous spans with the corresponding corrected text in ESC's outputs (the **Final Output** row).

Source Sentence	北京的空气太污染了，泛在北京的人一定要注意，别抽烟。
Annotation	北京的空气 <s1>太污染了，泛在北</s1> 京的人一定要注意，别抽烟。
Correction	<s1> 污染太严重了，在北</s1>
Final Output	北京的空气污染太严重了，在北京的人一定要注意，别抽烟。
Source Sentence	因为几乎的人们还没感到污染对自己的直接的影响。
Annotation	因 <s1>为几乎的人们还</s1> 没感到污染对自己的直接的影响。
Correction	<s1>为人们几乎</s1>
Final Output	因为为人们几乎还没感到污染对自己的直接的影响。
Source Sentence	列车、汽车，飞机等人类科技发展的结果也重大问题。
Annotation	列车、汽车，飞机等人类科技发展的结果 <s1>也重大</s1> 问题。
Correction	<s1>也存在重大</s1>
Final Output	列车、汽车，飞机等人类科技发展的结果也存在重大问题
Source Sentence	中国，悠久的历史，灿烂的文化，真是在历史上最难忘的国家。
Annotation	中国，悠久的历史，灿烂的文化，真 <s1>是在历</s1> 史上最难忘的国家。
Correction	<s1>是历</s1>
Final Output	中国，悠久的历史，灿烂的文化，真是历史上最难忘的国家。
Source Sentence	以找到这些稳定的工作，我们有读书的必要
Annotation	<s1>以找</s1> 到这些稳定的工作，我们有读书的必要。
Correction	<s1>为了找</s1>
Final Output	为了找到这些稳定的工作，我们有读书的必要

Table 10: Examples of our ESD & ESC approach in Chinese for GEC.