

An Alignment-Agnostic Model for Chinese Text Error Correction

Liying Zheng

liying.zheng@outlook.com

Yue Deng

dengyue606@gmail.com

Weishun Song

songweishun474@gmail.com

Liang Xu

xlpaul@126.com

Jing Xiao

xiaojing661@pingan.com.cn

Abstract

This paper investigates how to correct Chinese text errors with types of mistaken, missing and redundant characters, which are common for Chinese native speakers. Most existing models based on detect-correct framework can correct mistaken characters, but cannot handle missing or redundant characters due to inconsistency between model inputs and outputs. Although Seq2Seq-based or sequence tagging methods provide solutions to the three error types and achieved relatively good results in English context, they do not perform well in Chinese context according to our experiments. In our work, we propose a novel alignment-agnostic detect-correct framework that can handle both text aligned and non-aligned situations and can serve as a cold start model when no annotation data are provided. Experimental results on three datasets demonstrate that our method is effective and achieves a better performance than most recent published models.

1 Introduction

Chinese text error correction plays an important role in many NLP related scenarios (Martins and Silva, 2004; Afli et al., 2016; Wang et al., 2018; Burstein and Chodorow, 1999). For native Chinese speakers, common errors include mistaken characters, missing characters, and redundant characters. Mistaken characters refer to wrong characters needed to be replaced. Missing characters mean a lack of characters needed to be inserted into the identified position. Redundant characters mean useless or repeated characters needed to be deleted. Corrections for mistaken characters will not change the sentence length while corrections for the other two types will do. If texts only contain mistaken errors, we call it a text-aligned situation; if there exist missing or redundant errors, we call it a text non-aligned situation.

For text-aligned situation, many approaches apply the detect-correct framework, which is to detect the positions of wrong characters first and then correct them (Hong et al., 2019; Zhang et al., 2020; Cheng et al., 2020). Despite of competitive performance of such methods, they cannot deal with text non-aligned situation with missing and redundant errors. For text non-aligned situations, the reversed order error or complex structural change with multiple errors are not in our scope, first because we target to cover common mistakes made by Chinese native speakers, which are different from foreign Chinese learners in Chinese error correction(GEC) (Wang et al., 2020; Qiu and Qu, 2019) task, second because the mentioned complex errors are beyond our model settings. The two mainstream model schemes for text non-aligned situation are Seq2Seq-based and sequence tagging-based. The former is inspired by machine translation, which sets wrong sentences as input and correct sentences as output (Zhao et al., 2019; Kaneko et al., 2020; Chollampatt et al., 2019; Zhao and Wang, 2020; Lichtarge et al., 2019; Ge et al., 2018; Junczys-Dowmunt et al., 2018). Such approaches require a large number of training data and may generate uncontrollable results (Kiyono et al., 2019; Koehn and Knowles, 2017). The latter takes wrong sentences as input and modification operations of each token as output (Awasthi et al., 2019; Malmi et al., 2019; Omelianchuk et al., 2020). However, as Chinese language has more than 20,000 characters that can generate many combinations of token operations, it is difficult for sequence tagging models to cover all combinations and generate results with high coverage rates.

To address the above issues, we propose an alignment-agnostic detect-correct model, which can not only handle text non-aligned errors compared to the current detect-correct methods, but also can relieve the problem of huge value search space leading to uncontrollable or low covered

results of Seq2Seq or Sequence tagging based methods. We conduct experiments to compare our alignment-agnostic model with other models on three datasets: CGED 2020, SIGHAN 2015, SIGHAN-synthesize. Experimental results show that our model performs better than other models.

The contributions of our work include (1) proposal of a novel detect-correct architecture for Chinese text error correction, (2) empirical verification of the effectiveness of the alignment-agnostic model, (3) easy reproduction and fast adaptation to practical scenario with limited annotation data.

2 Our Approach

2.1 Problem Description

Chinese text error correction can be formalized as follows. Given a sequence of n characters $X = (x_1, x_2, x_3, \dots, x_n)$, the goal is to transform it into an m -character sequence $Y = (y_1, y_2, y_3, \dots, y_m)$, where n and m can be equal or not. The task can be viewed as a sequence transformation problem with a mapping function $f : X \rightarrow Y$

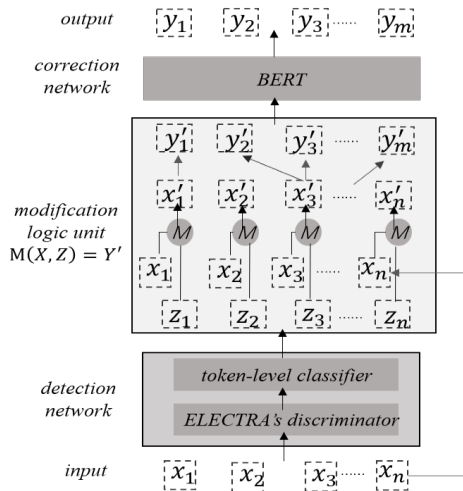


Figure 1: Architecture of the alignment-agnostic model

2.2 Model

As illustrated in Figure 1, the basic structure of our model includes a detection network evolved from ELECTRA discriminator (Clark et al., 2020) and a correction network based on BERT MLM (Devlin et al., 2019). The two networks are connected through a modification logic unit and are trained separately. The detection network locates the errors and identifies error types. The modification logic unit handles where and how to correct. Finally the correction network focuses on detailed correction.

The **detection network** is composed of an ELECTRA discriminator and a token-level error type classifier. The architecture of ELECTRA discriminator has been described in Clark et al. (2020). Here we modify the original classifier, and define the new token-level classifier with four categories, namely $label_{keep}$, $label_{mistaken}$, $label_{missing}$, $label_{redundant}$. $label_{keep}$ means the character is correct and should not change. $label_{mistaken}$ indicates the character is mistaken and needs to be replaced. $label_{missing}$ denotes we should insert characters before the current character. $label_{redundant}$ means the character is useless and needs to be deleted. We get the label probability of each token with the 4-class token-level classifier:

$$P_{i,label}(z_i = k|X) = softmax(w^T h^D(X)) \quad (1)$$

Where $P_{i,label}(z_i = k|X)$ denotes the conditional probability of character x_i being tagged with the label k , $h^D(X)$ is the last hidden state of ELECTRA discriminator and k is in label sets $[label_{keep}, label_{mistaken}, label_{missing}, label_{redundant}]$. The loss function of the detection network is:

$$Loss_{detect} = - \sum_{i=1}^n \log p_{i,label} \quad (2)$$

The **modification logic unit**, denoted by $M(X, Z)$, rewrites the input sequence X according to detection network's output Z :

$$if z_i = \begin{cases} label_{keep}, & x'_i = x_i \\ label_{mistaken}, & x'_i = [MASK] \\ label_{missing}, & x'_i = [MASK] x_i \\ label_{redundant}, & x'_i = '' \end{cases} \quad (3)$$

Based on the above formula, we get a new sequence $X' = (x'_1, x'_2, x'_3, \dots, x'_n)$. For each token with empty characters "", we delete it directly from the sequence X' , For each token with "[MASK]" x'_i , we reformulate it as two characters and obtain the final modified sequence $Y' = (y'_1, y'_2, y'_3, \dots, y'_m)$, whose length might be different from X' .

The **correction network** is BERT. We do the prediction for positions with the "[MASK]" symbol on the sequence Y' .

Test Set	Method	Detection			Correction		
		Prec.	Rec.	F_1 .	Prec.	Rec.	F_1 .
SIGHAN 2015	Hybrid (Wang et al., 2018)	56.6	69.4	62.3	-	-	57.1
	FASpell (Hong et al., 2019)	67.6	60	63.5	66.6	59.1	62.6
	Confusionset (Wang et al., 2018)	66.8	73.1	69.8	71.5	59.5	64.9
	Soft-Masked BERT(2020)	73.7	73.2	73.5	66.7	66.2	66.4
	our model(with a smaller training set)	79.1	64.0	71.3	72.2	60.6	68.2
	SpellGCN (Cheng et al., 2020)	74.8	80.7	77.7	72.1	77.7	75.9
	our model(with a larger training set)	87.5	68.6	76.9	87.0	65.2	74.6

Table 1: Performances of Different Methods on SIGHAN 2015. We trained our model on two datasets respectively.

3 Experiments

3.1 Datasets and Metrics

Chinese text error correction tasks mainly have two public datasets: the benchmark of SIGHAN 2015 (Tseng et al., 2015) which only contains text-aligned data and the competition of CGED 2020 (Rao et al., 2020) which contains text non-aligned data. In order to better verify our models’ effectiveness on text non-aligned scenario, we synthesized some non-aligned data based on SIGHAN 2015 dataset. Next, we will introduce how to utilize the three datasets.

For SIGHAN 2015 dataset, in order to keep accordance with other models in comparison, we incorporated SIGHAN 2013 and 2014 datasets in the training phase, as well as the SIGHAN 2013 confusion set. The test set contains 1100 passages and the train set contains 8738 passages. To ensure comparability, we also trained another model on a considerably larger train set to be consistent with SpellGCN’ (Cheng et al., 2020), which has 281379 passages in train set. We used the evaluation tool provided by SIGHAN, with metrics of precision (Prec.), recall(Rec.) and F_1 , all are based on sentence level.

CGED 2020 dataset is comprised of foreign Chinese learners’ writing, and contains an additional error type besides the three types mentioned above, which is the reversed order. As this type happens less frequently in native Chinese writing scenario, and is also beyond the scope of our model setting, we remove 575 relevant samples from a total sample of 2586, and get 846 training samples and 1165 testing samples. In consequence, we redo experiments with published models instead of comparing directly with the published benchmarks of other systems due to the inconsistency of test set.

To better verify our model’s effectiveness on

text non-aligned scenario, we synthesized some non-aligned data based on SIGHAN 2015 dataset (SIGHAN-synthesized). For mistaken characters error type, we kept the original errors unchanged. For missing characters error type, we randomly selected 50% samples and deleted one character from each of them. For redundant characters error type, we randomly selected 50% samples and inserted characters in each of them through four ways. (1) We inserted repeated characters in 35% of the selected samples. (2) We inserted confusing characters in 30% of the selected samples. (3) We inserted characters from high-frequency words in 30% of the selected samples. (4) We also inserted random characters in 5% of the selected samples.

For CGED 2020 dataset and SIGHAN-synthesized dataset, we adopted the M^2 score (Dahlmeier and Ng, 2012) and ER-RANT (Bryant et al., 2017) to evaluate models’ performance, which are two commonly used evaluation tools for text non-aligned situations.

3.2 Experiment Settings

The pre-trained ELECTRA discriminator model and BERT model adopted in our experiments are all from <https://github.com/huggingface/transformers>. We use the large-size ELECTRA and the base-size BERT. We train detection network and correction network on the three datasets respectively by Adam optimizer with default hyperparameters. All experiments are conducted on 2 GPUs (Nvidia Tesla P100).

For SIGHAN 2015, since it only contains one error type, we kept the default binary classifier of ELECTRA discriminator during finetuning. We applied two methods to retrain BERT. One is an unsupervised method by continue pretraining BERT with its original MLM objective. The other is a supervised method by masking mistaken characters

Test Set	Method	M_2 (Correction)			ERRANT(Correction)		
		Prec.	Rec.	$F_{0.5}$.	Prec.	Rec.	$F_{0.5}$.
CGED 2020	Copy-augmented(2019)	4.62	0.8	2.36	3.51	0.56	1.7
	Lasertagger(2019)	14.99	3.48	9.02	12.95	2.61	7.22
	PIE(2019)	22.3	10	17.9	17.1	6.6	13
	our model	29.71	22.03	27.77	24.8	17.56	22.91
SIGHAN-synthesized	Copy-augmented(2019)	38.44	8.03	21.87	38.31	7.8	21.5
	Lasertagger(2019)	51.29	43.21	49.44	50.14	39.99	47.72
	PIE(2019)	54.1	47.6	52.6	52	42.6	49.8
	our model	59.3	62.2	59.8	56.9	57.8	57

Table 2: The M^2 score and ERRANT score of Different Methods on CGED 2020 and SIGHAN-synthesized.

Method	SIGHAN 2015	CGED 2020	SIGHAN-synthesized
	F_1	$F_{0.5}$	$F_{0.5}$
ELECTRA+BERT	38.7	-	-
Finetune ELECTRA+BERT	66.2	22.91	54.7
Finetune ELECTRA+Finetune BERT	68.2	22.54	54.4
Finetune ELECTRA+Pretrain BERT	42	22.6	57

Table 3: Ablation study of alignment-agnostic model on three datasets.

and predicting them.

For CGED 2020 and SIGHAN-synthesized datasets, we added a 4-class classifier to recognize error types on ELECTRA discriminator’s last hidden layer and finetune it. We applied the same methods as in SIGHAN 2015 to retrain BERT.

3.3 Results

Table 1 shows the results on SIGHAN 2015 dataset. The first 5 lines implies that our method outperforms the method Soft-Masked BERT (Zhang et al., 2020) by 1.8% on F_1 score in correction phrase. With a larger train set, our model achieved higher F_1 score in both detection and correction phases. Compared with the previous SOTA method SpellGCN (Cheng et al., 2020), our model showed higher precision and comparable F_1 score.

Table 2 shows the results in comparison on CGED 2020 dataset and SIGHAN-synthesized dataset. Our model performs the best on correction level, exceeding the second best model by 9.87% on CGED 2020 and 7.2% on SIGHAN-synthesized dataset with $F_{0.5}M^2$ score. Since Copy-augmented (Zhao et al., 2019), as a Seq2Seq model, requires a large size of training data to get an acceptable result, it underperforms Lasertagger (Malmi et al., 2019) and PIE (Awasthi et al., 2019) models on both two datasets with a small training sample size. As analyzed before, sequence tagging models like Lasertagger and PIE do not

work well on Chinese language due to huge value search space.

3.4 Ablation Study

We carried out ablation study of our model on the three datasets. Table 3 shows the results on correction level. For SIGHAN 2015, finetuning ELECTRA can bring in a great improvement of 27.5% on F_1 score, while finetuning BERT only generates a relatively small rise of 2% on F_1 score and continue pretraining BERT leads to a decrease of 24.2% on F_1 score. A possible reason is that finetuning can incorporate confusion sets knowledge about similar characters easy to be mistaken, while unsupervised pretraining may destroy the original learned words distribution when training data largely differs from the original ones. Besides, our model achieves 38.7% on F_1 score with no training data and thus can work as a good baseline in cold start conditions. For CGED 2020 and SIGHAN-synthesized datasets, the two ways of retraining BERT didn’t improve much. Compared with the results of other SOTA models, the modification and finetuning of ELECTRA is the most effective part.

4 Conclusion

We proposed a new detect-correct model for Chinese text error correction. It can handle both text-aligned and non-aligned situations, and can serve

as a good baseline even in cold start situations. Experimental results on three datasets show that our model performs better than existing methods. Furthermore, it can be easily reproduced and achieve good results even with a small training data size, which is key to rapid application in the industry.

References

- Haithem Afli, Zhengwei Qiu, Andy Way, and Páiraic Sheridan. 2016. Using smt for ocr error correction of historical texts. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC'16)*, pages 962–966.
- Abhijeet Awasthi, Sunita Sarawagi, Rasna Goyal, Sabyasachi Ghosh, and Vihari Piratla. 2019. [Parallel iterative edit models for local sequence transduction](#). 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)*, pages 4260–4270, Hong Kong, China. Association for Computational Linguistics.
- Christopher Bryant, Mariano Felice, and Edward Briscoe. 2017. Automatic annotation and evaluation of error types for grammatical error correction. Association for Computational Linguistics.
- Jill Burstein and Martin Chodorow. 1999. Automated essay scoring for nonnative english speakers. In *Computer mediated language assessment and evaluation in natural language processing*.
- Xingyi Cheng, Weidi Xu, Kunlong Chen, Shaohua Jiang, Feng Wang, Taifeng Wang, Wei Chu, and Yuan Qi. 2020. [Spellgcn: Incorporating phonological and visual similarities into language models for chinese spelling check](#). *arXiv preprint arXiv:2004.14166*.
- Shamil Chollampatt, Weiqi Wang, and Hwee Tou Ng. 2019. [Cross-sentence grammatical error correction](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 435–445, Florence, Italy. Association for Computational Linguistics.
- Kevin Clark, Minh-Thang Luong, Quoc V Le, and Christopher D Manning. 2020. [Electra: Pre-training text encoders as discriminators rather than generators](#). *arXiv preprint arXiv:2003.10555*.
- Daniel Dahlmeier and Hwee Tou Ng. 2012. Better evaluation for grammatical error correction. In *Proceedings of the 2012 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 568–572.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [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)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Tao Ge, Furu Wei, and Ming Zhou. 2018. Reaching human-level performance in automatic grammatical error correction: An empirical study. *arXiv preprint arXiv:1807.01270*.
- Yuzhong Hong, Xianguo Yu, Neng He, Nan Liu, and Junhui Liu. 2019. [FASpell: A fast, adaptable, simple, powerful Chinese spell checker based on DAE-decoder paradigm](#). In *Proceedings of the 5th Workshop on Noisy User-generated Text (W-NUT 2019)*, pages 160–169, Hong Kong, China. Association for Computational Linguistics.
- Marcin Junczys-Dowmunt, Roman Grundkiewicz, Shubha Guha, and Kenneth Heafield. 2018. Approaching neural grammatical error correction as a low-resource machine translation task. *arXiv preprint arXiv:1804.05940*.
- 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](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4248–4254, Online. Association for Computational Linguistics.
- Shun Kiyono, Jun Suzuki, Masato Mita, Tomoya Mizumoto, and Kentaro Inui. 2019. [An empirical study of incorporating pseudo data into grammatical error correction](#). 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)*, pages 1236–1242, Hong Kong, China. Association for Computational Linguistics.
- Philipp Koehn and Rebecca Knowles. 2017. Six challenges for neural machine translation. *arXiv preprint arXiv:1706.03872*.
- Jared Lichtarge, Chris Alberti, Shankar Kumar, Noam Shazeer, Niki Parmar, and Simon Tong. 2019. [Corpora generation for grammatical error correction](#). pages 3291–3301.
- Eric Malmi, Sebastian Krause, Sascha Rothe, Daniil Mirylenka, and Aliaksei Severyn. 2019. [Encode, tag, realize: High-precision text editing](#). 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)*, pages 5054–5065, Hong Kong, China. Association for Computational Linguistics.

- Bruno Martins and Mário J Silva. 2004. [Spelling correction for search engine queries](#). In *International Conference on Natural Language Processing (in Spain)*, pages 372–383. Springer.
- Kostiantyn Omelianchuk, Vitaliy Atrasevych, Artem Chernodub, and Oleksandr Skurzhanskyi. 2020. [GECToR – grammatical error correction: Tag, not rewrite](#). In *Proceedings of the Fifteenth Workshop on Innovative Use of NLP for Building Educational Applications*, pages 163–170. Association for Computational Linguistics.
- Zhaoquan Qiu and Youli Qu. 2019. [A two-stage model for chinese grammatical error correction](#). *IEEE Access*, 7:146772–146777.
- Gaoqi Rao, Erhong Yang, and Baolin Zhang. 2020. Overview of nlptea-2020 shared task for chinese grammatical error diagnosis. In *Proceedings of the 6th Workshop on Natural Language Processing Techniques for Educational Applications*, pages 25–35.
- Yuen-Hsien Tseng, Lung-Hao Lee, Li-Ping Chang, and Hsin-Hsi Chen. 2015. Introduction to sighthan 2015 bake-off for chinese spelling check. In *Proceedings of the Eighth SIGHAN Workshop on Chinese Language Processing*, pages 32–37.
- Dingmin Wang, Yan Song, Jing Li, Jialong Han, and Haisong Zhang. 2018. [A hybrid approach to automatic corpus generation for Chinese spelling check](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2517–2527, Brussels, Belgium. Association for Computational Linguistics.
- Hongfei Wang, Michiki Kurosawa, Satoru Katsumata, and Mamoru Komachi. 2020. [Chinese grammatical correction using bert-based pre-trained model](#).
- Shaohua Zhang, Haoran Huang, Jicong Liu, and Hang Li. 2020. [Spelling error correction with soft-masked BERT](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 882–890, Online. Association for Computational Linguistics.
- 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](#). 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)*, pages 156–165, Minneapolis, Minnesota. Association for Computational Linguistics.
- Zewei Zhao and Houfeng Wang. 2020. Maskgec: Improving neural grammatical error correction via dynamic masking. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 1226–1233.