

Exploration and Exploitation: Two Ways to Improve Chinese Spelling Correction Models

Chong Li[†], Cenyuan Zhang[†], Xiaoqing Zheng^{*}, Xuanjing Huang

School of Computer Science, Fudan University, Shanghai, China

Shanghai Key Laboratory of Intelligent Information Processing

{chongli17, cenyuanzhang17, zhengxq, xjhuang}@fudan.edu.cn

Abstract

A sequence-to-sequence learning with neural networks has empirically proven to be an effective framework for Chinese Spelling Correction (CSC), which takes a sentence with some spelling errors as input and outputs the corrected one. However, CSC models may fail to correct spelling errors covered by the confusion sets, and also will encounter unseen ones. We propose a method, which continually identifies the weak spots of a model to generate more valuable training instances, and apply a task-specific pre-training strategy to enhance the model. The generated adversarial examples are gradually added to the training set. Experimental results show that such an adversarial training method combined with the pre-training strategy can improve both the generalization and robustness of multiple CSC models across three different datasets, achieving state-of-the-art performance for CSC task.¹

1 Introduction

Chinese Spelling Correction (CSC) aims to detect and correct spelling mistakes in Chinese texts. Many Chinese characters are visually or phonologically similar, while their semantic meaning may differ greatly. Spelling errors are usually caused by careless writing, automatic speech recognition, and optical character recognition systems. The CSC task has received steady attention over the past two decades (Chang, 1995; Xin et al., 2014; Wang et al., 2018; Hong et al., 2019). Unlike English, Chinese texts are written without using whitespace to delimit words, and it is hard to identify whether and which characters are misspelled without the information of word boundaries. The context information should be taken into account to reconstruct

the word boundaries when correcting spelling mistakes, which makes CSC a long-standing challenge for Chinese NLP community.

Many early CSC systems follow the same recipe with minor variations, adopting a three-step strategy: detect the positions of spelling errors; generate candidate characters for these positions; and select a most appropriate one from the candidates to replace the misspelling (Yeh et al., 2013; Yu and Li, 2014; Zhang et al., 2015; Wang et al., 2019). Recently, a sequence-to-sequence (seq2seq) learning framework with neural networks has empirically proven to be effective for CSC, which transforms a sentence with errors to the corrected one (Zhang et al., 2020; Cheng et al., 2020b).

However, even if training a CSC model with the seq2seq framework normally requires a huge amount of high-quality training data, it is still unreasonable to assume that all possible spelling errors have been covered by the *confusion* sets (i.e. a set of characters and their visually or phonologically similar characters which can be potentially confused) extracted from the training samples. New spelling errors occur everyday. A good CSC model should be able to exploit what it has already seen in the training instances in order to achieve reasonable performance on easy spelling mistakes, but it can also explore in order to generalize well to possible unseen misspellings.

In this study, we would like to pursue both the *exploration* (unknown misspellings) and *exploitation* (the spelling errors covered by the confusion sets) when training the CSC models. To encourage a model to explore unknown cases, we propose a character substitution-based method to pre-train the model. The training data generator chooses about 25% of the character positions at random for prediction. If a character is chosen, we replace it with the character randomly selected from its confusion set (90% of the time) or a random character (10%

[†]These authors contributed equally to this work.

^{*}Corresponding Author

¹The source codes are available at <https://github.com/FDChongli/TwoWaysToImproveCSC>.

of the time). Then, the model is asked to predict the original character.

Because of the combination of spelling errors and various contexts in which they occur, even though the confusion sets are given and fixed, models may still fail to correct characters that are replaced by any character from its confusion set. To better exploit what the models has experienced during the training phase, we generate more valuable training data via adversarial attack (i.e. tricking models to make false prediction by adding imperceptible perturbation to the input (Szegedy et al., 2014)), targeting the weak spots of the models, which can improve both the quality of training data for fine-tuning the CSC models and their robustness against adversarial attacks. Inspired by adversarial attack and defense in NLP (Jia and Liang, 2017; Zhao et al., 2018; Cheng et al., 2020a; Wang and Zheng, 2020), we propose a simple but efficient method for adversarial example generation: we first identify the most vulnerable characters with the lowest generation probabilities estimated by a pre-trained model, and replace them with characters from their confusion sets to create the adversarial examples.

Once the adversarial examples are obtained, they can be merged with the original clean data to train the CSC models. The examples generated by our method are more valuable than those already existed in the training set because they are generated towards to the weak spots of the current models. Through extensive experimentation, we show that such adversarial examples can improve both generalization and robustness of CSC models. If a model pre-trained with our proposed character substitution-based method is further fine-tuned by adversarial training, its robustness can be improved about 3.9% while without suffering too much loss (less than 1.1%) on the clean data.

2 Method

2.1 Problem Definition

Chinese Spelling Correction aims to identify incorrectly used characters in Chinese texts and giving its correct version. Given an input Chinese sentence $X = \{x_1, \dots, x_n\}$ consisting of n characters, which may contain some spelling errors, the model takes X as input and outputs an output sentence $Y = \{y_1, \dots, y_n\}$, where all the incorrect characters are expected to be corrected. This task can be formulated as a conditional generation problem by

modeling and maximizing the conditional probability of $P(Y|X)$.

2.2 Base Models

We use vanilla BERT (Devlin et al., 2019) and two recently proposed BERT-based models (Cheng et al., 2020b; Zhang et al., 2020) as our base models. When applying BERT to the CSC task, the input is a sentence with spelling errors, and the output representations are fed into an output layer to predict target tokens. We tie the input and output embedding layer, and all the parameters are fine-tuned using task-specific corpora. Soft-Masked BERT (Zhang et al., 2020) uses a Bi-GRU network to detect errors, and applies a BERT-based network to correct errors. SpellGCN (Cheng et al., 2020b) utilizes visual and phonological similarity knowledge through a specialized graph convolutional network and substitutes parameters of the output layer of BERT with the final output of it.

These models achieved state-of-the-art or close to state-of-the-art performance on the CSC task. However, we found that their performance and robustness could be further improved through pre-training and adversarial training, which help models explore unseen spelling errors and exploit weak points of themselves.

2.3 Pre-training Method

We collected unlabeled sentences from Wikipedia and Weibo corpora (Shang et al., 2015), covering both formal and informal Chinese texts. Training example pairs are generated by substituting characters in clean sentences, and models are trained to predict the original character. According to Chen et al. (2011), a sentence contains no more than two spelling errors on average, so we select and replace 25% characters in a sentence. The chosen Chinese character will be substituted by a character randomly selected from its confusion set (90% of the time) or a random Chinese character (10% of the time). The latter helps models to explore unknown misspellings not covered by the confusion sets.

2.4 Adversarial Example Generation and Adversarial Training

To efficiently identify and alleviate the weak spots of trained CSC models, we designed an adversarial attack algorithm for CSC tasks, which replaces the tokens in a sentence with spelling mistakes.

The adversarial examples generation algorithm in this paper can be divided into two main steps:

(1) determine the vulnerable tokens to change (2) replace them with the spelling mistakes that most likely to occur in the contexts (Algorithm 1).

For the i -th position of input sentence X , the positional score s_i can be obtained by the logit output o_i as follows:

$$s_i = o_i^{y_i} - o_i^{m_i} (o_i^{m_i} = \max\{o_i^r, r \neq y_i\}) \quad (1)$$

where o_i^r denotes the logit output of character r in the i -th position, and y_i denotes the i -th character of ground truth sentence Y . The lower the positional score, the less confident the model is in predicting the position. Attacking this position makes the model output more likely to change. Once the positional score of each character in the input sentence is calculated, we sort these positions in ascending order according to the positional scores. This process can reduce the substitutions and maintain the original semantics as much as possible.

Once a vulnerable position is determined, the token at that position is replaced with one of its phonologically or visually similar characters. Confusion set D contains a set of visually or phonologically similar characters. In order to fool the target CSC model while maintaining the context, the character with the highest logit output in the confusion set is used as a replacement.

Given a sentence in training sets, its adversarial examples are generated by substituting a few characters based on the algorithm mentioned above. Adversarial training was conducted with these examples, improving the robustness of CSC models by alleviating their weak spots, and exploiting knowledge about easy spelling mistakes from confusion sets to help models generalize better.

3 Experiments

3.1 Datasets

Statistics of the datasets used are shown in Table 1.

Pretraining data We generated a large corpus by a character substitution-based method. Models were first pre-trained on these nine million sentence pairs, and then fine-tuned using the training data mentioned below.

Training data The training data contained three human-annotated training datasets, SIGHAN 2013 (Wu et al., 2013), SIGHAN 2014 (Yu et al., 2014), and SIGHAN 2015 (Tseng et al., 2015). We also utilized an automatically generated dataset (Wang et al., 2018).

Algorithm 1 Adversarial Attack Algorithm

Input:

$X = \{x_1, x_2, \dots, x_n\}$, input Chinese sentence;
 $Y = \{y_1, y_2, \dots, y_n\}$, the corresponding ground truth;
 λ , proportion of characters can be changed;
 f , a target CSC model;
 D , a confusion set created based on visually or phonologically similar characters;

Output:

$\hat{X} = \{\hat{x}_1, \hat{x}_2, \dots, \hat{x}_n\}$, adversarial example;

```

1:  $\hat{X} \leftarrow X$ 
2: if  $f(X) \neq Y$  then
3:   return  $\hat{X}$ 
4: else
5:    $num \leftarrow 0$ 
6:   while  $f(\hat{X}) = Y$  &&  $num \leq \lambda \cdot n$  do
7:      $O = \{o_1, o_2, \dots, o_n\} \leftarrow$  Logit output of  $f(\hat{X})$ 
8:      $P = \{p_1, p_2, \dots, p_k\} \leftarrow$  Sort the position  $p_i$  in ascending order based on  $s_{p_i}$  ( $1 \leq p_i \leq n$  and  $y_{p_i}$  is a Chinese character)
9:     for each  $i \in [1, k]$  do
10:      if  $\hat{x}_{p_i} \neq y_{p_i}$  then
11:        continue
12:      end if
13:       $\hat{x}_{p_i} \leftarrow m_{p_i}$ , where  $m_{p_i} \in D(x_{p_i})$  and  $p_i^{m_{p_i}} = \max\{p_i^r, p_i^r \in D(x_{p_i})\}$ 
14:      break
15:    end for
16:     $num \leftarrow num + 1$ 
17:  end while
18: end if
19: return  $\hat{X}$ 

```

Table 1: Statistics information on the used data resources. A subset of the Wikipedia corpus and Weibo corpus, denoted by Wikipedia* and Weibo* respectively, was sampled from the entire corpus.

Pre-Training Data	#Line	Avg. Length	
Wikipedia*	4,531,007	40.2	
Weibo*	4,770,015	16.3	
Training Data	#Line	Avg. Length	#Errors
(Wang et al., 2018)	271,329	42.6	381,962
SIGHAN 2013	350	49.3	339
SIGHAN 2014	3,437	49.6	5,136
SIGHAN 2015	2,339	31.3	3,048
Test Data	#Line	Avg. Length	#Errors
SIGHAN 2013	1,000	74.3	1,221
SIGHAN 2014	1,062	50.0	771
SIGHAN 2015	1,100	30.7	705

Test data Models' performance in detection and correction stage was evaluated in sentence level on three benchmark datasets, in the metrics of F1 scores (detection and correction). Characters in these datasets were transferred into simplified Chinese characters using OpenCC². We revised the processed datasets for one simplified Chinese character may correspond to multiple traditional Chinese characters.

²<https://github.com/BYVoid/OpenCC>

Table 2: Performance of three models trained with the proposed pretraining strategy and adversarial training method. “CLEAN” stands for the testing results on the clean data, and “ATTACK” denotes the F1 scores under test-time attacks. “DET” and “COR” denote the F1 scores of detection and correction. The F1 scores were increased 4.1% on average by our pre-training method across the various models on the different datasets. Models’ robustness was also improved about 3.9% while without suffering too much loss (less than 1.1%) on the clean data.

Model	SIGHAN-2013				SIGHAN-2014				SIGHAN-2015			
	CLEAN		ATTACK		CLEAN		ATTACK		CLEAN		ATTACK	
	DET	COR	DET	COR	DET	COR	DET	COR	DET	COR	DET	COR
BERT	82.9	82.1	33.6	15.8	66.8	65.0	41.7	19.0	76.3	74.4	25.1	13.7
+ Pre-trained for CSC	84.9	84.4	48.5	29.6	70.4	68.6	51.4	32.4	79.8	78.0	39.0	26.9
+ Adversarial training	84.0	83.5	50.8	31.3	68.4	66.8	54.9	38.0	80.0	78.2	45.9	36.0
SpellGCN	80.8	80.0	25.6	22.6	64.8	63.6	29.0	24.3	73.6	71.5	18.8	17.4
+ Pre-trained for CSC	84.6	84.0	28.8	25.8	67.3	66.4	35.4	27.1	79.6	77.7	26.2	25.2
+ Adversarial training	83.4	82.6	30.2	26.0	66.4	65.4	35.9	29.5	79.6	77.8	28.2	25.2
Soft-masked BERT	80.6	79.1	27.7	4.0	62.2	59.6	29.8	7.1	72.4	69.6	15.5	5.3
+ Pre-trained for CSC	84.9	84.2	27.3	6.0	67.2	65.6	30.7	8.6	77.2	74.5	22.2	6.5
+ Adversarial training	84.1	83.3	32.5	8.1	65.0	62.7	40.5	13.4	76.2	73.8	30.3	11.4

3.2 Models and Hyper-parameter Settings

For BERT and Soft-Masked BERT, we used the BERT model pre-trained on Chinese text provided by transformers³ and fine-tuned it. Adam optimizer was used and the learning rate was $2e-5$, except when adversarial training on SIGHAN 13 dataset, which was $1e-5$. We followed Zhang et al. (2020) to set our hyper-parameters. The size of the hidden state in Bi-GRU in Soft-Masked BERT was 256.

Similarly, we followed the hyper-parameters settings of SpellGCN (Cheng et al., 2020b) except the batch size. Batch size was reduced to eight due to GPU memory. The BERT model used in SpellGCN was provided by the repository of BERT⁴.

We conducted adversarial training on base models gained through pre-training and fine-tuning. The threshold λ was tuned on the validation set for each dataset. The number of sentence pairs directly used for training was twice that that used to generate adversarial examples.

3.3 Results and Analysis

As shown in Table 2, through pre-training particularly designed for CSC, the models achieve better results on three benchmark datasets. The average improvement of correction F1 score was 4.3% over base CSC models, which proves that our pre-training method has significant contribution to improving the model. Notably, BERT achieves state-of-the-art results on three datasets through our method.

³<https://github.com/huggingface/transformers>

⁴<https://github.com/google-research/bert>




Figure 1: Trade-off between generalization and robustness. The blue and orange lines respectively denote the average F1 scores of BERT on the SIGHAN-2015 data set and the adversarial examples generated ($\lambda = 0.05$).

Figure 1 shows the trade-off between generalization and robustness during adversarial training. As the threshold increases, the robustness of BERT also increases with a slight performance decrease on clean dataset (less than 0.7%).

The experiments of the models under adversarial attacks were conducted with the base, pre-trained and adversarially trained models ($\lambda = 0.02$). We found that CSC models are vulnerable to adversarial examples as expected. The average drop in F1 score of three base models was 51.6%. Under the attacks, the F1 scores of adversarially trained model decreased less (44.1%), which indicates the adversarial training can substantially improve the robustness of CSC models. Compared with other models, BERT is more robust against adversarial attack (-41.2%). The reason for the more serious

robustness issues of other models may be related to the modules added to BERT, which increases the number of parameters, therefore it is more likely to overfit on the CSC data set.

4 Conclusion

In this paper, we have described a character substitution-based method to create large pseudo data to pre-train the models by encouraging them to explore unseen misspellings. We also proposed a data augmentation method for training the CSC models by continually adding the adversarial examples, particularly generated to alleviate the weak spot of the current model, to the training set. By the proposed pre-training strategy and adversarial training method, we can pursue both the exploration and exploitation when training the CSC models. Experimental results demonstrate that the CSC models trained with the data augmented by these pseudo data and adversarial examples can substantially be improved in both generalization and robustness.

Acknowledgements

The authors would like to thank the anonymous reviewers for their valuable comments. This work was supported by National Key R&D Program of China (No. 2018YFC0830900), Shanghai Municipal Science and Technology Major Project (No. 2021SHZDZX0103), National Science Foundation of China (No. 62076068) and Zhangjiang Lab.

References

- Chao-Huang Chang. 1995. A new approach for automatic chinese spelling correction. In *Proceedings of Natural Language Processing Pacific Rim Symposium*, volume 95, pages 278–283. Citeseer.
- Yong-Zhi Chen, Shih-Hung Wu, Ping-Che Yang, Tsun Ku, and Gwo-Dong Chen. 2011. Improve the detection of improperly used chinese characters in students’ essays with error model. *International Journal of Continuing Engineering Education and Life Long Learning*, 21(1):103–116.
- Minhao Cheng, Jinfeng Yi, Pin-Yu Chen, Huan Zhang, and Cho-Jui Hsieh. 2020a. [Seq2sick: Evaluating the robustness of sequence-to-sequence models with adversarial examples](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(04):3601–3608.
- Xingyi Cheng, Weidi Xu, Kunlong Chen, Shaohua Jiang, Feng Wang, Taifeng Wang, Wei Chu, and Yuan Qi. 2020b. [SpellGCN: Incorporating phonological and visual similarities into language models for Chinese spelling check](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 871–881, Online. Association for Computational Linguistics.
- 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.
- 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.
- Robin Jia and Percy Liang. 2017. [Adversarial examples for evaluating reading comprehension systems](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2021–2031, Copenhagen, Denmark. Association for Computational Linguistics.
- Lifeng Shang, Zhengdong Lu, and Hang Li. 2015. [Neural responding machine for short-text conversation](#). 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)*, pages 1577–1586, Beijing, China. Association for Computational Linguistics.
- Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. 2014. [Intriguing properties of neural networks](#). In *International Conference on Learning Representations*.
- Yuen-Hsien Tseng, Lung-Hao Lee, Li-Ping Chang, and Hsin-Hsi Chen. 2015. [Introduction to SIGHAN 2015 bake-off for Chinese spelling check](#). In *Proceedings of the Eighth SIGHAN Workshop on Chinese Language Processing*, pages 32–37, Beijing, China. Association for Computational Linguistics.
- 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 (EMNLP)*, Brussels, Belgium.
- Dingmin Wang, Yi Tay, and Li Zhong. 2019. [Confusionset-guided pointer networks for Chinese spelling check](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5780–5785, Florence, Italy. Association for Computational Linguistics.

- Lihao Wang and Xiaoqing Zheng. 2020. [Improving grammatical error correction models with purpose-built adversarial examples](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2858–2869, Online. Association for Computational Linguistics.
- Shih-Hung Wu, Chao-Lin Liu, and Lung-Hao Lee. 2013. [Chinese spelling check evaluation at SIGHAN bake-off 2013](#). In *Proceedings of the Seventh SIGHAN Workshop on Chinese Language Processing*, pages 35–42, Nagoya, Japan. Asian Federation of Natural Language Processing.
- Yang Xin, Hai Zhao, Yuzhu Wang, and Zhongye Jia. 2014. [An improved graph model for Chinese spell checking](#). In *Proceedings of The Third CIPS-SIGHAN Joint Conference on Chinese Language Processing*, pages 157–166, Wuhan, China. Association for Computational Linguistics.
- Jui-Feng Yeh, Sheng-Feng Li, Mei-Rong Wu, Wen-Yi Chen, and Mao-Chuan Su. 2013. [Chinese word spelling correction based on n-gram ranked inverted index list](#). In *Proceedings of the Seventh SIGHAN Workshop on Chinese Language Processing*, pages 43–48, Nagoya, Japan. Asian Federation of Natural Language Processing.
- Junjie Yu and Zhenghua Li. 2014. [Chinese spelling error detection and correction based on language model, pronunciation, and shape](#). In *Proceedings of The Third CIPS-SIGHAN Joint Conference on Chinese Language Processing*, pages 220–223, Wuhan, China. Association for Computational Linguistics.
- Liang-Chih Yu, Lung-Hao Lee, Yuen-Hsien Tseng, and Hsin-Hsi Chen. 2014. [Overview of SIGHAN 2014 bake-off for Chinese spelling check](#). In *Proceedings of The Third CIPS-SIGHAN Joint Conference on Chinese Language Processing*, pages 126–132, Wuhan, China. Association for Computational Linguistics.
- 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.
- Shuiyuan Zhang, Jinhua Xiong, Jianpeng Hou, Qiao Zhang, and Xueqi Cheng. 2015. [HANSpeller++: A unified framework for Chinese spelling correction](#). In *Proceedings of the Eighth SIGHAN Workshop on Chinese Language Processing*, pages 38–45, Beijing, China. Association for Computational Linguistics.
- Zhengli Zhao, Dheeru Dua, and Sameer Singh. 2018. [Generating natural adversarial examples](#). In *International Conference on Learning Representations*.