

Mask the Correct Tokens: An Embarrassingly Simple Approach for Error Correction

Kai Shen^{1*}, Yichong Leng^{2*}, Xu Tan^{3†}

Siliang Tang^{1†}, Yuan Zhang⁴, Wenjie Liu⁴, Edward Lin⁴

¹Zhejiang University, ²University of Science and Technology of China

³Microsoft Research Asia, ⁴Microsoft Azure Speech

¹{shenkai, siliang}@zju.edu.cn, ²lyc123go@mail.ustc.edu.cn

^{3,4}{xuta, yuanz, liwenjie, edlin}@microsoft.com

Abstract

Text error correction aims to correct the errors in text sequences such as those typed by humans or generated by speech recognition models. Previous error correction methods usually take the source (incorrect) sentence as encoder input and generate the target (correct) sentence through the decoder. Since the error rate of the incorrect sentence is usually low (e.g., 10%), the correction model can only learn to correct on limited error tokens but trivially copy on most tokens (correct tokens), which harms the effective training of error correction. In this paper, we argue that the correct tokens should be better utilized to facilitate effective training and then propose a simple yet effective masking strategy to achieve this goal. Specifically, we randomly mask out a part of the correct tokens in the source sentence and let the model learn to not only correct the original error tokens but also predict the masked tokens based on their context information. Our method enjoys several advantages: 1) it alleviates trivial copy; 2) it leverages effective training signals from correct tokens; 3) it is a plug-and-play module and can be applied to different models and tasks. Experiments on spelling error correction and speech recognition error correction on Mandarin datasets and grammar error correction on English datasets with both autoregressive and non-autoregressive generation models show that our method improves the correction accuracy consistently¹.

1 Introduction

Text error correction techniques have been widely used to correct the errors in text, such as grammar errors (Rothe et al., 2021; Grundkiewicz et al.,

2019; Omelianchuk et al., 2020) or spelling errors (Cheng et al., 2020; Zhang et al., 2020) and automatic speech recognition (ASR) errors (Cucu et al., 2013; D’Haro and Banchs, 2016; Anantaram et al., 2018; Zhu et al., 2021), etc. Text error correction is a challenging task because in realistic ASR transcripts or human-typed documents, it is common that the errors are sparse (for example, the error rate of an ASR system is usually below 10%) (Leng et al., 2021b). To handle the problem, previous works aim to improve the accuracy of correction models by using error detector-corrector architecture and focusing on two directions: 1) promoting the accuracy of error detection (Zhang et al., 2020), 2) enhancing the capability of correction (Cheng et al., 2020; Leng et al., 2021b,a). Unfortunately, they all fall in the short that the conventional correction models always copy the correct tokens directly, which is potentially harmful to the learning of the error correction process.

The learning diagram that directly copies the correct tokens can be harmful in two aspects. Firstly, it affects the learning of the correction procedure of the error tokens. Conceptually, the correction diagram can be divided into two main streams: 1) copy the correct token, 2) correct the error token. Intuitively, the former is much easier and contains less language understanding knowledge. The unbalance of trivial copying task (about 90% tokens) and hard correction task (about 10% tokens) is potentially harmful to the training of the error correction model. Secondly, it severely underuses valuable training data and consequently lead to low training efficiency. To be more specific, since the errors in input text are sparse, most of the tokens in input text are correct and the correction model only learns a trivial copy mapping from input to output, which means that the model can hardly learn knowledge such as language understanding from correct tokens other than copying. To the best of our knowledge, there is no previous work proposed

* Both authors contributed equally. This work was conducted when the first and second authors were interns in Microsoft.

† Corresponding author

¹The code for our model is provided for research purposes: <https://github.com/microsoft/NeuralSpeech>.

to better leverage the correct tokens.

In this paper, we propose a simple but effective framework MaskCorrect to alleviate the heavy copy phenomenon and improve the utilization of correct tokens in training data. In detail, during the training of the correction model, we randomly perturb a proportion of correct tokens to a special token <mask> in the source sentences and train the correction model to predict the correct tokens from <mask> token. With the help of our method, the correction model can benefit from several aspects: 1) it learns the language understanding knowledge on correct tokens instead of a trivial copy mapping, and thus alleviating the trivial copy phenomenon; 2) since the <mask> token does not contain explicit clues about its corresponding token, our method will encourage the model to generate output tokens based more on context information, and thus it alleviates the low training efficiency problem.

We conduct a series of experiments on two tasks: 1) Spelling Error Correction and 2) ASR Error Correction on Mandarin datasets with both auto-regressive and non-autoregressive methods. Furthermore, we also conduct experiments on the Grammar Error Correction task on English datasets in Appendix D.1. Experimental results show that our MaskCorrect can consistently boost the conventional correction methods in both tasks. Especially, we outperform the state-of-the-art spelling error correction method by a large margin.

In addition, it is worth noting that the method proposed in this paper is a general framework and can be potentially extended in situations where there exists lots of trivial copy mapping in training data, such as text summarization and question generation.

2 Related Works

In this section, we briefly review the spelling/ASR error correction methods and then discuss the copy phenomenon in closely related NLP tasks.

Spelling/ASR Error Correction The mainstream model on spelling error correction leverage a detector-corrector framework (Cheng et al., 2020; Zhang et al., 2020; Guo et al., 2021), which consists of an error detector to detect error tokens and an error corrector to correct the tokens detected as error tokens. A series of works are proposed to improve the correction accuracy with the help of pretraining methods such as BERT (Zhang et al., 2020) or pronunciation of word (Zhang et al., 2021). As for

the ASR error correction, the early-stage correction model for ASR leverages the auto-regressive translation model to correct the errors (Mani et al., 2020; Liao et al., 2020; Wang et al., 2020; Zhu et al., 2021). Recently, non-autoregressive correction models are proposed to correct errors in a fast and accurate manner (Leng et al., 2021b,a; Du et al., 2022). However, all the methods in both tasks focus on the learning of correcting error tokens, and the models can only learn a trivial copy mapping for tokens without error. Furthermore, they also neglect the low training efficiency because of the heavy copy phenomenon.

The Copy Phenomenon in NLP tasks The copy phenomenon and the problems it causes have been noticed in the related NLP tasks. In the abstractive text, the copy (pointer) mechanism has been widely integrated to address the issue of frequently copying words from the source context to the desired summary (See et al., 2017; Wu et al., 2020). However, a series of works raise potential shorts. Wang et al. (2019); Ji et al. (2020) finds that merely copying parts of the original text tends to plain and uncreative summary. Wei et al. (2019) proposes a novel regularization method to solve the dilemma that traditional copy mechanisms cannot work properly when the target tokens are not consistent with the source tokens.

When it comes to language pretraining, the copy phenomenon attracts lots of attention from researchers. In the original MLM training, it is a common practice (Devlin et al., 2019; Liu et al., 2019; He et al., 2020; Li et al., 2022b) to only mask up to 15% of original tokens and recover them from the corrupted context while the remaining other tokens are not trained. ELECTRA (Clark et al., 2020) claims that the MLM training only learns from the 15% of the tokens, leading to the severe substantial compute cost and low training efficiency. Instead of MLM, they introduce a new pretraining diagram that some tokens are firstly randomly replaced with alternative tokens and the model is trained with the replaced token detection task where it learns to distinguish real input tokens with plausible generated tokens. Wettig et al. (2022) finds that masking more than 15% of tokens can bring benefits. With a large mask ratio, the model can learn from more cases and thus boosting the training efficiency.

Although the copy phenomenon is still common in the correction, compared to the pretraining which aims to learn the general and robust contex-

tual knowledge, the correction model has a much cleaner purpose that recovering the incorrect tokens mistaken by humans. In this paper, we propose to perturb correct tokens to utilize these correction tokens to help the model learn knowledge about context understanding.

3 Method

3.1 Preliminary Discussion

In this section, we first conduct some preliminary experiments to qualitatively discuss the impact of the copy phenomenon. In addition, we discuss a simple solution to alleviate the proposed copy problem. Without loss of generality, we conduct preliminary experiments on FastCorrect (Leng et al., 2021b) which is a state of the art (SOTA) baseline in ASR Error Correction (we will introduce the task and FastCorrect in Sec. 4.2).

Add Additional "copy" Tokens. Based on the assumption that heavy copy is harmful to the correction, one natural question is: what will happen if we make the "copy" more frequently?

To answer this question, we gradually increase the "copy" tokens in the training set while keeping the number of the error tokens unchanged. In detail, we first find out all the samples denoted as \mathcal{C} in the paired training data whose input sentence is equal to the target sentence. Then we randomly select q (denoted as the "copy" ratio q) samples from collection \mathcal{C} and add them to the paired training set. Finally, we train the FastCorrect on the synthetic training set. We vary the "copy" ratio $q \in \{0.1, 0.2, 0.3, 0.4, 0.5\}$. We report the word error rate (WER) metric on the AISHELL-1 (Bu et al., 2017)'s test set in Figure 1.

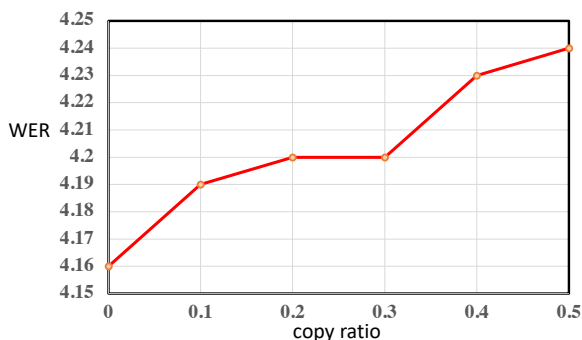


Figure 1: The performance of FastCorrect with different ratios of training "copy" ratio in test set of AISHELL-1. The WER is in %. Note: the smaller WER is, the better performance the correction model has.

From the results, we can find that as the "copy" ratio increases, the performance of FastCorrect will decrease gradually. Although FastCorrect is trained with the same amount of error samples, the model is paying more attention to the simply-copied correct samples. This phenomenon illustrates that the heavy copy phenomenon is harmful to the correction, which proves the correctness of our motivation.

Mask on Target Loss. To alleviate the heavy copy problem, one direct but natural idea is to mask off the gradients of some correct tokens during the gradient back-propagation. We implement this idea by simply masking off a specific number of correct tokens in the loss calculation. The results are shown in Figure 2.




Figure 2: The performance of FastCorrect with different mask ratios of target correct tokens in the test set of AISHELL-1. The WER is in %. Note: the smaller WER is, the better performance the correction model has.

From the results, surprisingly, we can find that the WER results increase when we gradually mask more correct tokens. This experiment indicates that although the correct tokens are copied as a phenomenon, they cannot be easily dropped during the loss calculation. We have several insights: 1) we cannot alleviate the "copy" by preventing the loss in back-propagation. Every token plays an important role in the context and the overall distribution will be influenced if we simply prevent back-propagating gradient of some tokens; 2) Although the correct tokens are copied as a result, they still contain information in some way and cannot be ignored easily.

Motivated by this experiment, in the MaskCorrect framework, we choose to mask the correct tokens in the source sentence and predict the corresponding token. There are several advantages: 1) every token participates in the correction loss

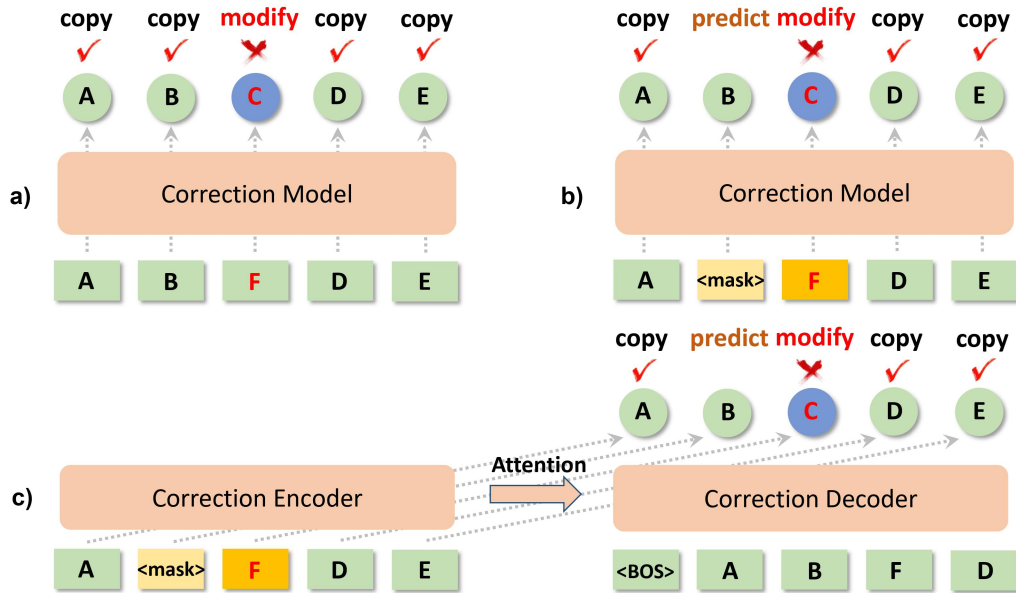


Figure 3: The conventional correction framework (a) and our proposed MaskCorrect in both non-autoregressive (b) and autoregressive (c) framework.

back-propagation, which does not affect the overall distribution; 2) by masking off the correct tokens in the source sentence, the model is encouraged to predict the token only by the context.

3.2 MaskCorrect

In this subsection, we introduce our method for addressing the heavy copy of correct tokens in the training data of correction models.

As shown in Figure 3 (a), the conventional correction models focus on the correction of error tokens (i.e., token “F” in Figure 3) and simply learn a copy mapping on the rest of the tokens that have no error (i.e., token “A”, “B”, “D” and “E” in Figure 3). If the effective training samples in training data are the tokens with error only, then the valuable training data is heavily underused since only a small proportion of tokens in training data (typically about 10%) contains errors.

We encourage the model to leverage correct tokens by introducing perturbation, as shown in Figure 3 (b) in non-autoregressive and (c) in autoregressive way. Specifically, we random perturb $p\%$ input tokens to a special token `<mask>` before the input tokens are fed into the correction model and keep the target tokens untouched. As a result, the model needs to predict the masked token based on its context information (i.e., predicting token “B” based on its context information in Figure 1 (b) and (c)). A better understanding of context information will also benefit the correction of error tokens since

the model needs to leverage context information to detect which token has an error and how to correct this error token.

We would like to highlight that the perturbation is only conducted on the input tokens, and thus our method can be used on any error correction tasks with arbitrary model architecture or loss function designed for certain error correction tasks. Our perturbation helps the model better utilize correction tokens by a mask-and-predict task.

4 Experiments

In this section, we evaluate our method on two tasks, namely *ASR Error Correction* and *Spelling Error Correction* on Mandarin datasets. We first report the accuracy of the baselines and our methods to demonstrate the effectiveness of our method on each task. Then we conduct ablation experiments to comprehensively analyze the effect of the mask mechanism. Furthermore, we also conduct experiments on *Grammar Error Correction* on English corpus in Appendix D.1 to demonstrate the effectiveness.

4.1 Spelling Error Correction

In this section, we evaluate the effectiveness of our method on spelling error correction task.

4.1.1 Datasets

We conduct the experiments on two publicly available SIGHAN datasets: SIGHAN 2014 (Yu and Li,

Test Sets	Model	Detection Level				Correction Level			
		Acc.	P.	R.	F1	Acc.	P.	R.	F1
SIGHAN14	FASPELL	-	61.00	53.50	57.00	-	59.40	52.00	55.40
	SpellGCN	69.83	69.20	65.20	67.14	68.92	70.10	60.94	65.20
	SoftMask	69.49	72.69	60.38	65.97	68.36	71.90	58.08	64.26
	+ MaskCorrect	71.37	74.11	63.85	68.60	70.72	73.70	62.50	67.64
	BERT	68.27	69.10	63.65	66.27	66.85	68.10	60.77	64.23
	+ MaskCorrect	71.94	73.82	66.15	69.78	70.90	73.19	64.04	68.31
SIGHAN15	FASPELL	-	67.60	60.00	63.50	-	66.60	59.10	62.60
	SpellGCN	78.92	80.80	75.70	78.17	78.10	80.10	72.70	76.22
	SoftMask	77.36	79.96	72.14	75.85	76.45	79.54	70.30	74.63
	+ MaskCorrect	78.45	80.32	74.54	77.32	77.09	79.71	71.77	75.53
	BERT	78.82	82.12	72.88	77.22	77.64	81.62	70.48	75.64
	+ MaskCorrect	80.91	81.56	79.15	80.34	79.00	80.79	75.28	77.94

Table 1: The performance of the baselines and our method for Spelling Error Correction. We report the accuracy (Acc.), precision (P.), recall (R.) and F1 scores. We report the FASPELL’s performance in (Cheng et al., 2020). Best results are in bold. All metrics are in %.

2014) (SIGHAN14) and SIGHAN 2015 (Tseng et al., 2015) (SIGHAN15) from Chinese Spell Check Bake-offs. Following Cheng et al. (2020); Guo et al. (2021), we adopt the standard split of the training and testing data of SIGHAN, as well as the same data preprocessing which converts the traditional Chinese to simple Chinese by OpenCC.

From the previous works (Cheng et al., 2020; Guo et al., 2021), the SIGHAN is too small to achieve acceptable results, we follow their common practice and collect 5 million unpaired texts from the internet to pretrain the model. The detailed procedure to add noises can be found in Appendix B.1. For the finetune split, we include the training corpus of SIGHAN14 and SIGHAN15. In addition, we also include the 271K samples generated by automatic method (Wang et al., 2018) as the finetune data². The detailed data statistics are shown in Appendix A.1.

We use the pretrained Bert tokenizer provided by HuggingFace (Wolf et al., 2020) to tokenize the sentences. In addition, we filter out the sentences whose length are larger than 100 for the paired and unpaired training corpus.

4.1.2 Baseline Methods

FASPELL The FASPELL (Hong et al., 2019) proposes a new correction paradigm that consists of

a denoising autoencoder and a decoder. Instead of using an inflexible and insufficient confusion dictionary, it utilizes the similarity metric to select the candidates in the decoder.

SpellGCN the SpellGCN (Cheng et al., 2020) incorporates the phonological and visual similarity knowledge into language modeling. They propose to build the similarity graph for each word token in the view of the pronunciation and word shape.

BERT The BERT (Devlin et al., 2018) is a commonly used baseline for spelling error correction. It simply employs the pretrained BERT as a sentence encoder network. We take the encoded sentence features and feed them to network with multiple linear layers and ReLU activation to predict the targets.

SoftMask BERT This method (Zhang et al., 2020) is a BERT-based spelling error correction baseline, which consists of a GRU-based error detection module and a BERT-based error correction module. A soft-masking gate softly fuses the <mask> token’s embedding with the initial word embedding by the probability of being an error token predicted by the detection module. Then the fused embeddings are fed to the correction module.

4.1.3 Evaluation Metrics

We report the sentence level accuracy, precision, recall, and F1 scores as the evaluation metrics, which are commonly used in the spelling error correction

¹<http://ir.itc.ntnu.edu.tw/lre/sighan7csc.html>

²<https://github.com/wdimmy/Automatic-Corpus-Generation/blob/master/corpus/train.sgml>

	Mask Ratio	Detection Level				Correction Level			
		Acc.(%)	P.(%)	R.(%)	F1(%)	Acc.(%)	P.(%)	R.(%)	F1(%)
SIGHAN14	0%	68.27	69.10	63.65	66.27	66.85	68.10	60.77	64.23
	10%	65.82	66.05	64.23	65.49	65.82	66.05	62.12	64.02
	15%	71.85	73.87	65.77	69.58	70.80	73.29	63.85	68.24
	20%	71.94	73.82	66.15	69.78	70.90	73.19	64.04	68.31
	25%	72.60	76.08	64.23	69.66	71.56	75.47	62.12	68.14
	30%	68.36	68.85	64.62	66.67	67.80	68.46	63.46	65.87
	40%	66.29	66.01	64.23	65.11	64.78	64.90	61.15	62.97
	50%	67.61	68.26	63.27	65.67	66.67	67.58	61.35	64.31
SIGHAN15	0%	78.82	82.12	72.88	77.22	77.64	81.62	70.48	75.64
	10%	76.64	76.84	75.28	76.05	75.27	76.16	72.51	74.29
	15%	78.27	78.32	77.31	77.81	77.36	77.90	75.46	76.66
	20%	80.91	81.56	79.15	80.34	79.00	80.79	75.28	77.94
	25%	78.36	79.92	74.91	77.33	77.09	79.35	72.32	75.68
	30%	76.36	75.73	76.57	76.15	74.82	74.95	73.43	74.18
	40%	77.82	76.90	78.60	77.74	76.55	76.30	76.01	76.16
	50%	76.82	77.54	74.54	76.01	75.18	76.74	71.22	73.88

Table 2: The ablation experiments on the mask ratio hyper-parameter p . The best metrics are in bold. All the metrics are in % and Acc.-Accuracy, P.-Precision, R.-Recall.

tasks. All the metrics are computed by the official evaluation tools¹.

4.1.4 Implementation Details

We choose the *BERT* and *SoftMask BERT* as the correction backbone and evaluate the effectiveness of our method. The detailed training configuration can be found in Appendix C.1. For both BERT and SoftMask BERT, given the input sentence with m tokens, we first find out all the correct tokens by comparing the input sentence with the ground truth sentence. Secondly, we randomly select the tokens from the correct token set with ratio p without replacement. For each selected token, it will be replaced with <mask> token with a probability of m , or will be replaced with a randomly chosen other token with a probability of n , otherwise it will not be changed with a probability of $1 - m - n$. We apply the MaskCorrect on both unpaired and paired training corpus with the same mask hyper-parameter p, m, n . However, since it is inevitable that there is no such <mask> token during the inference stage, we have to finetune model with the paired data additionally for 3 epochs to cope with the training-inference mismatch problem to achieve better performance.

4.1.5 Results Analysis

In this section, we compare our method with the baseline models on the two datasets for Spelling Error Correction. We apply the proposed MaskCorrect method on two baselines, namely *BERT + MaskCorrect* and *SoftMask + MaskCorrect*. The results are shown in Table 1.

First, we compare all the methods and we find the variant *BERT + MaskCorrect* outperforms all other baselines in both SIGHAN14 and SIGHAN15 datasets. Specifically, we compare two baselines *BERT* and *SoftMask* with the corresponding MaskCorrect variants. With the MaskCorrect, 1) the *SoftMask* obtains 2.63% and 3.02% F1 score gains on detection and correction levels in SIGHAN14 and obtains 1.47% and 0.9% gains on SIGHAN15, respectively; 2) the *BERT* obtains 3.51% and 4.07% F1 gains and 3.12% and 2.30%, respectively. It demonstrates that by randomly perturb a certain proportion of correct tokens into <mask> tokens, the correction models can learn the semantic knowledge from the correct tokens, which enhances the ability of detection and correction the spelling errors.

4.1.6 Ablation Study

In this section, we further study the impact of the masking ratio (i.e., the hyper-parameter

¹<http://nlp.ee.ncu.edu.tw/resource/csc.html>

p). Without loss of generality, we choose the best baseline *BERT + MaskCorrect* as the framework and conduct experiments on both SIGHAN14 and SIGHAN15 with p varying in $\{0.1, 0.15, 0.2, 0.25, 0.3, 0.4, 0.5\}$. The results are shown in Table 2.

From the results, we have several observations: 1) when the mask ratio is small, there are no gains or even a slight performance drop. Because when the mask ratio is small, there is not enough knowledge learned from the correct tokens and the mask tokens may hinder the correction of error tokens. 2) when the mask ratio is large, the MaskCorrect cannot bring the expected benefits. Empirically, we find that under a large mask ratio, the cross-entropy loss is always large and cannot be well-optimized even with many epochs. This is because when we mask too much correct tokens, the language information loses too much and the language model cannot learn the correct knowledge from it. 3) the best choice of mask ratio is 0.2 in this task for the BERT correction model.

4.2 ASR Error Correction

In this section, we evaluate the effectiveness of our method on the ASR Error Correction task.

4.2.1 Dataset and ASR Model

We conduct the experiments on two Chinese datasets, the public dataset AISHELL-1 (Bu et al., 2017) and the inhouse dataset. Following the previous works (Zhu et al., 2021), we employ the ASR model to train on the training set. Then the ASR model is used to transcribe the ASR training split again to generate the transcripts. The transcripts along with the corresponding ground truth texts are used to train the ASR correction model. The validation and test splits for ASR correction are inferred by the above ASR model. The details of dataset statistics and ASR model can be found in Appendix A.2 and B.3, respectively.

However, for AISHELL-1 dataset, it’s too small to achieve good correction performance (Leng et al., 2021b,a). To address this issue, we follow the previous works’ (Leng et al., 2021b,a) practice to construct a much larger pseudo training corpus to pretrain the correction model. Details can be found in Appendix B.2. As for the inhouse dataset, since it’s large enough to achieve good results, we don’t apply unpaired data pretraining and directly train the correction model from scratch.

4.2.2 Baseline Methods and Evaluation Metric

FELIX FELIX (Mallinson et al., 2020) is a NAR algorithm for text edition. It aligns the mismatch tokens between the source and target sentences by two sub-tasks: 1) decide the subset of tokens in source sentences and their order in the target sentences by tagging operation and 2) insert tokens to infill the missing tokens in the target sentences.

LevT Levenshtein Transformer (LevT) (Gu et al., 2019) is a partially autoregressive seq2seq model devised for flexible sequence generation tasks. It implicitly predicts the insertion and deletion with multiple iterations in the decoder.

Autoregressive Correction Model (AR Transformer) For the autoregressive sequence generation model, we follow the standard Transformer (Vaswani et al., 2017) which takes the sentences with potential errors as input and outputs the corrected sentences autoregressively.

FastCorrect FastCorrect (Leng et al., 2021b) is a non-autoregressive ASR correction model based on edit distance. It employs edit distance to find the fine-grained mapping between error tokens and the corresponding correct tokens based on edit operation: insertion, deletion and substitution. We follow the default setting and implement it based on the official code¹.

As for the evaluation metrics, we follow the practice in previous works and adopt the word error rate (WER) and Word Error Reduction Rate (WERR) as the automatic evaluation metric. We use the official code in ESPnet to calculate it.

4.2.3 Implementation Details

We evaluate our method on two SOTA ASR correction model in both AR and NAR, namely *AR Transformer* and *FastCorrect*. Similar to the Spelling Error Correction 4.1.4, we first collect the correct tokens from inputs and randomly mask the tokens with ratio p and probability m and n . It is worth noting that we apply the MaskCorrect on both unpaired and paired training corpus for AISHELL-1 and only paired training corpus for inhouse dataset (since we do not use unpaired training corpus for the inhouse dataset as mentioned in Section 4.2.1). In addition, since there is no <mask> token during inference, we have to finetune on the original

¹<https://github.com/microsoft/NeuralSpeech/tree/master/FastCorrect>

Methods	AISHELL-1				Inhouse Corpus	
	Test Set		Dev Set		Test Set	
	WER ↓	WERR ↑	WER ↓	WERR ↑	WER ↓	WERR ↑
No Correction	4.83	-	4.46	-	8.99	-
<i>Non-autoregressive correction models</i>						
LevT(MIter=1)	4.73	2.07	4.37	2.02	8.86	1.45
LevT(MIter=3)	4.74	1.86	4.38	1.79	8.89	1.11
FELIX	4.63	4.14	4.26	4.48	8.74	2.78
FastCorrect	4.16	13.87	3.89	13.30	8.44	6.12
+ MaskCorrect	4.10	15.11	3.82	14.35	8.35	7.12
<i>Autoregressive correction models</i>						
AR Transformer	4.08	15.53	3.80	14.80	8.32	7.45
+ MaskCorrect	4.03	16.56	3.75	15.92	8.25	8.23

Table 3: The performance of the baseline models and our method for ASR Error Correction. Note that we only apply our method on the strongest non-autoregressive correction model. We report the WER and the WERR as the evaluation metrics. All metrics are in %. "MIter" is a hyper-parameter in LevT controlling max decoding iteration.

data without mask for 3 epochs. Please refer to Appendix C.2 for more details.

4.2.4 Results Analysis

In this section, we compare our method with the baseline models on the two datasets for ASR Error Correction. We apply the proposed MaskCorrect framework on two baselines: *FastCorrect* and *AR Transformer*, for NAR and AR correction diagrams, respectively. Note that we only apply our method to the strongest non-autoregressive correction model, *FastCorrect*. The results are shown in Table 3.

From the table, we can have the following observations. Firstly, the *AR + MaskCorrect* model achieves the best performance compared with all other baselines, which demonstrates the effectiveness of our method. Secondly, by comparing the *FastCorrect* and *AR Transformer* with the corresponding variants, we can find that the *MaskCorrect* can enhance both baselines on all datasets. Concretely, with the *MaskCorrect*, the *FastCorrect* can further reduce the WER (measured by WERR) of ASR model by 1.24% and 1% on the test sets of the AISHELL-1 and inhouse dataset, respectively. Meanwhile the *AR Transformer* can reduce further the WER by 1.03% and 0.78%, respectively. It confirms that the *MaskCorrect* is useful for both AR and NAR methods. By adding a perturb (i.e., masking off some correct tokens), the correction model can learn from the correct tokens, which is beneficial to the ASR error correction.

4.2.5 Ablation Study

In this section, we conduct an ablation study to verify the affect of the mask ratio p . We choose the best NAR baseline *FastCorrect + MaskCorrect* as the framework and apply different value of p varying in $\{0.1, 0.15, 0.2, 0.25, 0.3, 0.4, 0.5\}$ on the training set. The results are shown in Table 4.

From the results, we have several observations: 1) when the mask ratio falls in $[0, 0.2]$, the *MaskCorrect* will bring a positive impact to the *FastCorrect*, otherwise it will cause a slight performance drop. We think when the mask ratio is too large, the semantic information contained in the incomplete text (with mask tokens and error tokens) is missing too much. Thus, it is too difficult for *FastCorrect* to recover it. When the mask ratio is in a suitable range, *FastCorrect* will benefit from recovering the mask tokens. 2) the best choice of mask ratio is 0.15 for the *FastCorrect*.

Mask Ratio	AISHELL-1		Inhouse Corpus	
	WER	WERR	WER	WERR
0%	4.16	13.87	8.38	6.79
10%	4.14	15.75	8.35	7.12
15%	4.10	16.67	8.32	7.45
20%	4.16	15.30	8.34	7.23
25%	4.19	14.61	8.34	7.23
30%	4.2	14.38	8.38	6.79
40%	4.22	13.93	8.40	6.56
50%	4.24	13.47	8.41	6.45

Table 4: The ablation experiments on the mask ratio hyper-parameter p for ASR Error Correction. The best metrics are in bold. All the metrics are in %.

5 Conclusion

In this work, we proposed an effective and plug-and-play framework to address the trivial copy phenomenon, in which we randomly perturb a proportion of correct tokens into a special token <mask>. The model aims to predict the correct tokens from the special token. We apply this mechanism to both AR and NAR models for Spelling Error Correction tasks and the ASR Correction tasks. And experimental results show the effectiveness of the proposed method. Furthermore, our method can be easily extended to other tasks which suffer from heavy copy phenomenon.

6 Acknowledgements

This work has been supported in part by the Zhejiang NSF (LR21F020004), Chinese Knowledge Center of Engineering Science and Technology (CKCEST).

7 Limitation

In this work, we propose a framework MaskCorrect to address the heavily copy phenomenon in the spelling and ASR correction tasks. We partially solve the problem by a simple mask strategy with quantitative experiments. But, a more in-depth theoretical analysis of how to prevent the negative copy impact and how to adaptively mask the copied tokens is worth studying. Furthermore, despite the learning diagram we have explored, how to adapt the conventional model structure to cope with the heavily copied phenomenon is another future direction.

References

- C Anantaram, Amit Sangroya, Mrinal Rawat, and Aishwarya Chhabra. 2018. Repairing asr output by artificial development and ontology based learning. In *IJCAI*, pages 5799–5801.
- 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.
- Hui Bu, Jiayu Du, Xingyu Na, Bengu Wu, and Hao Zheng. 2017. Aishell-1: An open-source mandarin speech corpus and a speech recognition baseline. In *2017 20th conference of the oriental chapter of the international coordinating committee on speech databases and speech I/O systems and assessment (O-COCOSDA)*, pages 1–5. IEEE.
- 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. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 871–881, Online. 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*.
- Horia Cucu, Andi Buzo, Laurent Besacier, and Corneliu Burileanu. 2013. Statistical error correction methods for domain-specific asr systems. In *International Conference on Statistical Language and Speech Processing*.
- 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.
- 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*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In *NAACL*, pages 4171–4186.
- Jing Du, Shiliang Pu, Qinbo Dong, Chao Jin, Xin Qi, Dian Gu, Ru Wu, and Hongwei Zhou. 2022. Cross-modal asr post-processing system for error correction and utterance rejection. In *arXiv preprint arXiv:2201.03313*.
- Luis Fernando D’Haro and Rafael E Banchs. 2016. Automatic correction of asr outputs by using machine translation. In *INTERSPEECH*.
- Sylviane Granger. 2014. The computer learner corpus: a versatile new source of data for sla research. In *Learner English on computer*, pages 3–18. Routledge.
- 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. *Advances in Neural Information Processing Systems*, 32.
- Zhao Guo, Yuan Ni, Keqiang Wang, Wei Zhu, and Guotong Xie. 2021. Global attention decoder for chinese spelling error correction. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 1419–1428.
- Pengcheng He, Xiaodong Liu, Jianfeng Gao, and Weizhu Chen. 2020. Deberta: Decoding-enhanced bert with disentangled attention. *arXiv preprint arXiv:2006.03654*.
- 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.
- Changzhen Ji, Xin Zhou, Yating Zhang, Xiaozhong Liu, Changlong Sun, Conghui Zhu, and Tiejun Zhao. 2020. Cross copy network for dialogue generation. *arXiv preprint arXiv:2010.11539*.
- Taku Kudo and John Richardson. 2018. Sentencepiece: A simple and language independent subword tokenizer and detokenizer for neural text processing. *arXiv preprint arXiv:1808.06226*.
- Yichong Leng, Xu Tan, Rui Wang, Linchen Zhu, Jin Xu, Wenjie Liu, Linqun Liu, Tao Qin, Xiangyang Li, Edward Lin, et al. 2021a. Fastcorrect 2: Fast error correction on multiple candidates for automatic speech recognition. *arXiv preprint arXiv:2109.14420*.
- Yichong Leng, Xu Tan, Linchen Zhu, Jin Xu, Renqian Luo, Linqun Liu, Tao Qin, Xiangyang Li, Edward Lin, and Tie-Yan Liu. 2021b. Fastcorrect: Fast error correction with edit alignment for automatic speech recognition. *Advances in Neural Information Processing Systems*, 34:21708–21719.
- Jiquan Li, Junliang Guo, Yongxin Zhu, Xin Sheng, Deqiang Jiang, Bo Ren, and Linli Xu. 2022a. Sequence-to-action: Grammatical error correction with action guided sequence generation. *arXiv preprint arXiv:2205.10884*.
- Juncheng Li, Xin He, Longhui Wei, Long Qian, Linchao Zhu, Lingxi Xie, Yueting Zhuang, Qi Tian, and Siliang Tang. 2022b. Fine-grained semantically aligned vision-language pre-training. *Advances in Neural Information Processing Systems*.
- Junwei Liao, Sefik Emre Eskimez, Liyang Lu, Yu Shi, Ming Gong, Linjun Shou, Hong Qu, and Michael Zeng. 2020. Improving readability for automatic speech recognition transcription. In *arXiv preprint arXiv:2004.04438*.
- 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. *arXiv preprint arXiv:1907.11692*.
- Jonathan Mallinson, Aliaksei Severyn, Eric Malmi, and Guillermo Garrido. 2020. Felix: Flexible text editing through tagging and insertion. *arXiv preprint arXiv:2003.10687*.
- Anirudh Mani, Shruti Palaskar, Nimshi Venkat Meripo, Sandeep Konam, and Florian Metze. 2020. Asr error correction and domain adaptation using machine translation. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*.
- 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.
- 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.
- Hwee Tou Ng, Siew Mei Wu, Yuanbin Wu, Christian Hadiwinoto, and Joel Tetreault. 2013. **The CoNLL-2013 shared task on grammatical error correction**. In *Proceedings of the Seventeenth Conference on Computational Natural Language Learning: Shared Task*, pages 1–12, Sofia, Bulgaria. Association for Computational Linguistics.
- Kostiantyn Omelianchuk, Vitaliy Atrasevych, Artem Chernodub, and Oleksandr Skurzshanskyi. 2020. Gector—grammatical error correction: tag, not rewrite. *arXiv preprint arXiv:2005.12592*.
- Sascha Rothe, Jonathan Mallinson, Eric Malmi, Sebastian Krause, and Aliaksei Severyn. 2021. A simple recipe for multilingual grammatical error correction. *arXiv preprint arXiv:2106.03830*.
- Abigail See, Peter J Liu, and Christopher D Manning. 2017. Get to the point: Summarization with pointer-generator networks. *arXiv preprint arXiv:1704.04368*.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2015. Neural machine translation of rare words with subword units. *arXiv preprint arXiv:1508.07909*.
- 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.

- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- 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.
- Haoyu Wang, Shuyan Dong, Yue Liu, James Logan, Ashish Kumar Agrawal, and Yang Liu. 2020. Asr error correction with augmented transformer for entity retrieval. In *Interspeech*, pages 1550–1554.
- Wenbo Wang, Yang Gao, He-Yan Huang, and Yuxiang Zhou. 2019. Concept pointer network for abstractive summarization. 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 3076–3085.
- Bingzhen Wei, Xuancheng Ren, Yi Zhang, Xiaoyan Cai, Qi Su, and Xu Sun. 2019. Regularizing output distribution of abstractive chinese social media text summarization for improved semantic consistency. *ACM Transactions on Asian and Low-Resource Language Information Processing (TALLIP)*, 18(3):1–15.
- Alexander Wettig, Tianyu Gao, Zexuan Zhong, and Danqi Chen. 2022. Should you mask 15% in masked language modeling? *arXiv preprint arXiv:2202.08005*.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. 2020. *Transformers: State-of-the-art natural language processing*. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, 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.
- Yiquan Wu, Kun Kuang, Yating Zhang, Xiaozhong Liu, Changlong Sun, Jun Xiao, Yueting Zhuang, Luo Si, and Fei Wu. 2020. De-biased court’s view generation with causality. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 763–780.
- 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*, pages 180–189.
- 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.
- Ruiqing Zhang, Chao Pang, Chuanqiang Zhang, Shuo-huan Wang, Zhongjun He, Yu Sun, Hua Wu, and Haifeng Wang. 2021. Correcting Chinese spelling errors with phonetic pre-training. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 2250–2261.
- 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.
- Yu Zhang, Guoguo Chen, Dong Yu, Kaisheng Yao, Sanjeev Khudanpur, and James Glass. 2016. Highway long short-term memory rnns for distant speech recognition. In *2016 IEEE international conference on acoustics, speech and signal processing (ICASSP)*, pages 5755–5759. IEEE.
- Linchun Zhu, Wenjie Liu, Linqun Liu, and Edward Lin. 2021. Improving asr error correction using n-best hypotheses. In *IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*.

A Dataset Statistics

A.1 The Dataset Statistics for Spelling Error Correction

We provide the detailed dataset statistics for spelling error correction in Table 5.

Traning data	# Sent	Avg. Len
Unpaired Corpus	5 million	51
Aug.	271329	44.4
SIGHAN 14	6526	49.7
SIGHAN 15	3174	30.0
Paired Corpus Sum.	281029	44.3
Test Data	# Sent	Avg. Len
SIGHAN 14	1062	50.1
SIGHAN 15	1100	30.5

Table 5: The dataset statistics for Spelling Error Correction.

A.2 The dataset statistics for ASR Error Correction

The AISHELL-1 dataset is a Mandarin Speech corpus containing 178 hours of training audios, 10 hours of validation audios and 5 hours of test audios. And the inhouse dataset is a much larger dataset for the industrial Mandarin ASR system. It contains 75K hours of Mandarin speech data for training and 200 hours of speech for both validation and test.

B Dataset

B.1 Add Noise to unpaired data for Spelling Error Correction

We take the official confusion dictionary released by (Wu et al., 2013)¹, which is a collection of mapping from one Chinese character to its' widely mistaken characters. Following (Zhang et al., 2020), we randomly replace 15% of the characters in them with the mistaken characters by the confusion dictionary to artificially generate corresponding sentences with errors.

B.2 Add Noise to unpaired data for ASR Error Correction

We first crawl 400M unpaired sentences from the Internet. Secondly, we find out the insertion, deletion, and substitution error rate in the training set. Finally, for each unpaired sentence, we apply random insertion, deletion, and substitution with a homophone dictionary based on the error rate statistics.

B.3 The ASR Model for ASR Error Correction

As for the ASR model, we adopt a conformer with 12 layers encoder and 6 layers decoder, which is trained with cross-entropy loss and an auxiliary CTC loss on the encoder output. The hidden size of the conformer's encoder and decoder layers is 512. We use ESPnet¹ as the codebase. For the inhouse dataset, the ASR model is a hybrid model consisted of a latency-controlled BLSTM (Zhang et al., 2016) acoustic model(am) and a 5-gram language model which is trained on 436 billion tokens.

¹<https://github.com/espnet/espnet>

C Implementation Details

C.1 Implementation Details of Spelling Error Correction

For a fair comparison, we choose the same pre-trained BERT model for the above two methods provided by HuggingFace². For the BERT correction model, we use 2 layers of MLPs with softmax to transform the bert embeddings to the target token probability distribution. Another 2 layers of MLPs are used to predict the error probability as an assistant loss. For the SoftMask BERT correction model, we use a bidirectional GRU with hidden dimension 768 as the error detection network. The MLP layers' dimension is set to 1024. The learning rate is set to 1e-4. We pretrain on the unpaired data with 9 epochs and finetune on the paired data with 4 epochs. The pretrain batch size is 1024 and the finetune batch size is 128. The mask probability of m is 80%, and the randomly replace probability of n is 10%. We mask the correct tokens with probability $p = 20%$.

C.2 Implementation Details of ASR Error Correction

We use SentencePiece (Kudo and Richardson, 2018) to learn the subwords and apply it to all the texts above. The dictionary size is set to 40K. We train all correction models on 4 NVIDIA Tesla V100 GPUs, with a batch size of 12000 tokens. The mask probability of m is 80%, and the randomly replace probability of n is 10%. We mask the correct tokens with probability $p = 15%$.

D Experiments

D.1 Grammar Error Correction

In this section, we evaluate the effectiveness of our method on the Grammar Error Correction (GEC) task.

D.1.1 Datasets

We conduct the experiments on the datasets of BEA-2019 GEC shared task (Bryant et al., 2019), restricted track. Specifically, the training set is consisted by four parts: Lang-8 (Mizumoto et al., 2011), the FCE training set (Yannakoudakis et al., 2011), NUCLE (Dahlmeier et al., 2013), and W&I+LOCNESS (Granger, 2014; Bryant et al., 2019). We use the CoNLL-2013 (Ng et al., 2013) test set as our dev set and the CoNLL-2014 (Ng

²<https://huggingface.co/bert-base-chinese>

et al., 2014) test set as our test set. We follow the S2A’s (Li et al., 2022a) practice and preprocess and tokenize all the sentences by 32K Sentence-Piece (Kudo and Richardson, 2018) Byte Pair Encoding (BPE) (Sennrich et al., 2015).

We use the official MaxMatch (M^2) scorer (Dahlmeier and Ng, 2012) to evaluate the models in the English GEC tasks. Given a source sentence, and a ground-truth sentence sample, the M^2 scorer searches for the highest overlap between them and reports the precision, recall, and $F_{0.5}$.

D.1.2 Baseline Methods

Transformer-big The transformer (Vaswani et al., 2017) architecture has shown its success in the sequence-to-sequence fields. We simply cast the GEC task as a sequence-to-sequence problem, which takes the sentences with errors as inputs and predicts the correct sentences autoregressively. We follow the default setting of transformer-big and implement it based on the official codes¹.

Sequence-to-Action (S2A) The sequence-to-action (Li et al., 2022a) is a strong baseline based on sequence tagging, which is another diagram compared with seq2seq methods for GEC. The S2A method first predicts the token-level sequence actions and fuses the actions with the seq2seq framework to generate the final sentences.

D.1.3 Implementation Details

We evaluate our method on the transformer-big architecture. Similar to the previous tasks, we first collect the correct tokens from inputs and randomly mask the tokens with ratio 0.2 and probability $m = 0.8$ and $n = 0.1$. We finetune the model by 2 epochs.

D.1.4 Results Analysis

In this section, we compare our method with the baseline models for the GEC task. The results are shown in Table 6. From the results, we have several observations. Firstly, our method improves the transformer-big baseline for the English GEC task, which demonstrates the effectiveness of MaskCorrect. Secondly, with MaskCorrect, the transformer-big achieves state-of-art performance on all metrics.

¹<https://github.com/facebookresearch/fairseq>

Model	Precision	Recall	$F_{0.5}$
Transformer-big	64.9	26.6	50.4
Transformer-big + MaskCorrect	67.2	30.0	53.9
Sequence-to-Action	65.9	28.9	52.5

Table 6: The dataset statistics for Spelling Error Correction.