

Efficient Grammatical Error Correction Via Multi-Task Training and Optimized Training Schedule

Andrey Bout

Huawei Noah’s Ark Lab
bout.andrey@huawei.com

Alexander Podolskiy

Huawei Noah’s Ark Lab
podolskiy.alexander@huawei.com

Sergey Nikolenko

St. Petersburg Department of the
Steklov Institute of Mathematics
sergey@logic.pdmi.ras.ru

Irina Piontkovskaya

Huawei Noah’s Ark Lab
piontkovskaya.irina@huawei.com

Abstract

Progress in neural grammatical error correction (GEC) is hindered by the lack of annotated training data. Sufficient amounts of high-quality manually annotated data are not available, so recent research has relied on generating synthetic data, pretraining on it, and then fine-tuning on real datasets; performance gains have been achieved either by ensembling or by using huge pretrained models such as XXL-T5 as the backbone. In this work, we explore an orthogonal direction: how to use available data more efficiently. First, we propose auxiliary tasks that exploit the alignment between the original and corrected sentences, such as predicting a sequence of corrections. We formulate each task as a sequence-to-sequence problem and perform multi-task training. Second, we discover that the order of datasets used for training and even individual instances within a dataset may have important effects on the final performance, so we set out to find the best training schedule. Together, these two ideas lead to significant improvements, producing results that improve state of the art with much smaller models; in particular, we outperform the best models based on T5-XXL (11B parameters) with a BART-based model (400M parameters).

1 Introduction

Grammatical error correction (GEC) is an important problem setting with obvious applications that often require high-level understanding. Recent research has developed a common view of grammatical error correction (GEC) as monolingual text-to-text rewriting (Náplava and Straka, 2019; Grundkiewicz et al., 2019). GEC is often tackled with encoder-decoder architectures: the encoder receives an errorful sentence, and the decoder pro-

duces its corrected version. In line with mainstream NLP developments, Transformer-based architectures have proven to be beneficial for GEC, leading to recent breakthroughs on multiple benchmarks (Rothe et al., 2021; Tarnavskiy et al., 2022; Sun and Wang, 2022).

GEC is a challenging task that goes beyond simple spell-checking. The greatest difficulty comes from the complexity of grammar and punctuation rules, inconsistency in rules usage, contextual factors, and multiple correct answers. To solve GEC, a model should have a thorough understanding of grammar, punctuation, and syntax, and be aware of the ambiguity introduced by errors to correctly understand intended meaning.

Most recent works approach these challenges by providing new complex architectures, using larger models, or ensembling multiple models, e.g. Yuan et al. (2021); Rothe et al. (2021); Tarnavskiy et al. (2022); Zhang et al. (2022). However, practicality demands faster and smaller models, so in this work, we examine another direction to improve GEC: modifying the training pipeline. Every step is important: (i) *data preparation* includes generating synthetic data for pretraining and preprocessing available datasets; (ii) *pretraining* is done on synthetic data and/or large-scale datasets with supervised pretext tasks such as language modeling; (iii) *fine-tuning* on downstream datasets is also far from straightforward; at the very least, one has to choose the *order* of training on available datasets, which may significantly impact the results. Due to the scarcity of annotated data for fine-tuning, we are concerned with using it as efficiently as possible and introduce two novel approaches for this.

First, one can acquire additional information from parallel annotated corpora, e.g., an alignment

between two sentences that can be used to derive *operations* that would transform one into the other. One can decompose the editing process into elementary edits: insertion, deletion, replacement, and change of word order; the latter can also be viewed as a sequence of deletions and insertions. This information is exploited by sequence tagging models that show good results on the GEC task, see Omelianchuk et al. (2020, 2021); Tarnavskiy et al. (2022). A drawback of this approach lies in the need to specify a vocabulary of available operations, use several decoding steps to correct complex errors, and in the hardness/inability to make complex rephrasing. We utilize additional information by converting it into auxiliary tasks and formulating them as sequence-to-sequence problems. As such, the model learns separate “skills” required for the successful correction of grammatical errors.

Second, GEC-specific datasets have different quality and sources, e.g., different language proficiency levels. Some datasets are collected from on-line platforms where users correct each other (they are noisier); in other datasets, sentences are corrected by teachers or annotators (these are cleaner). Some datasets consist of essays written by native speakers; others, of essays written by non-native students. Thus, the distribution of errors may severely differ. It is tempting to remove “noisy” or out-of-distribution examples from training sets, but it seems that the model can learn even from such instances. We propose to use a training schedule for GEC datasets and show that it does matter how we order them. Also, we find that the order of sentences within the dataset matters as well; namely, we have found that placing sentences from the same block of text (e.g., the same essay) in the same batch is beneficial.

Our primary contributions are as follows. First, we introduce a multi-task pretraining approach and a fine-tuning schema that together yield an improvement of up to 3% in $F_{0.5}$ score compared to similar-sized state of the art models¹. Second, we show that our approach is able to outperform state of the art *large* models (T5-XL with 3B parameters and T5-XXL with 11B parameters) using a model with 400M parameters, reducing the computational load and ecological footprint. Third, our approach makes the model more robust, improving metrics on several datasets rather than trading them off

¹We plan to release the source code of our models upon acceptance.

Dataset	Sentences	% errorful	Stages
C4 _{200M}	~ 180M	99.4	I
PIE-synthetic	~ 9M	100.0	I
Lang-8	947 344	52.5	II
NUCLE	56 958	38.0	II
FCE	34 490	62.4	II
W&I+L	34 304	67.3	II, III
W&I+L dev	4 384	64.3	Dev
CoNLL test	1 312	71.9	Test
W&I+L test	4 477	N/A	Test

Table 1: Dataset statistics and training stages.

of each other (as usual). In what follows, Section 2 surveys related work, Section 3 introduces our approach, Section 4 shows evaluation results, an ablation study, and an analysis of our auxiliary tasks, Section 5 concludes the paper, and Section 6 discusses the limitations of our approach.

2 Related work

Neural approaches to grammatical error correction follow two main lines of research: (i) sequence tagging models and (ii) sequence-to-sequence models.

Synthetic data. Training GEC models is difficult due to the natural lack of suitable training data and possible erroneous corrections, so synthetic data becomes a crucial part of any GEC pipeline (Choe et al., 2019; Stahlberg and Kumar, 2021; Htut and Tetreault, 2019). It had been used for GEC even before the deep learning era that required larger datasets (Foster and Andersen, 2009; Brockett et al., 2006). Synthetic data generators can mimic common typos and grammatical errors but usually cannot capture the target error distribution found in real-life evaluation sets or standard benchmarks. Methods for synthetic data generation include character perturbations, dictionary- or edit-distance-based replacements, shuffling word order, rule-based suffix transformations, and more (Grundkiewicz et al., 2019; Awasthi et al., 2019a; Náplava and Straka, 2019; Rothe et al., 2021). An empirical study of how to generate and use the synthetic data was done in Kiyono et al. (2019).

Another line of research upsamples training data in existing datasets. Mita et al. (2020) train a GEC model on a natural noisy dataset and then use its outputs for source sentences to construct a less noisy parallel dataset; Zhang et al. (2019)

use sentence rewriting approaches; [Lichtarge et al. \(2020\)](#) apply delta-log-perplexity scoring to rank sentences according to the difference in perplexity between two base model checkpoints and use higher-scoring sentences for final fine-tuning.

Multi-stage fine-tuning. Due to data scarcity, training GEC models from scratch could be cumbersome. One of the options is to pre-train a model on some auxiliary task, e.g. [Choe et al. \(2019\)](#) proposed to initialize the GEC model with the pre-trained denoising autoencoder. Many GEC pipelines utilize pre-trained language models as backbone models for GEC; in particular, [Rothe et al. \(2021\)](#) used T5 ([Raffel et al., 2020](#)) while [Katsumata and Komachi \(2020\)](#) and [Sun and Wang \(2022\)](#) used BART ([Lewis et al., 2020](#)). Pre-trained language models are also beneficial for reranking output hypotheses generated with beam search ([Kaneko et al., 2019](#)). [Choe et al. \(2019\)](#), [Omelianchuk et al. \(2020\)](#) and [Tarnavskiy et al. \(2022\)](#) decompose the training process of a GEC model into several stages: (i) pre-training on an errorful synthetic dataset; (ii) fine-tuning on natural high-quality datasets that combine both errorful and error-free sentences. Each stage requires its own tuning of hyperparameters such as the number of training steps and the learning rate.

Multi-task learning. Several works aim to utilize additional information along with the standard parallel mapping. First, the grammatical error detection (GED) task can be extracted from GEC; [Yuan et al. \(2019\)](#) perform multi-task training with GED and GEC tasks and use GED features for reranking. A similar approach by [Fang et al. \(2020\)](#) trained a GED model separately and used it to filter edits. [Yuan et al. \(2021\)](#) separately pretrain a GED model and use its outputs as auxiliary inputs to fine-tune the encoder-decoder GEC model and rerank its outputs. [Zhang et al. \(2022\)](#) incorporate syntactic dependency information into the encoder.

Non-autoregressive decoding. Another line of research introduces non-autoregressive decoding to speed up models. [Awasthi et al. \(2019b\)](#) predict language-specific edits to be applied to the output sequence. Iterative refinement is also possible. [Gu et al. \(2019\)](#) non-autoregressively refine an output sequence using language-agnostic insertions and deletions. [Yakovlev et al. \(2023\)](#) decompose the inference stage into permutation and decoding. First, a permutation network repositions the tokens of an input sequence with possible deletions and inser-

tions. Second, the intermediate sequence is passed to a decoder network that iteratively fills in inserted placeholders.

GPT-3.5 and GPT-4. The recent success of GPT-based models for a wide variety of tasks has led to several parallel works that compare how well these models fare in grammatical error correction. [Fang et al. \(2023\)](#) show that ChatGPT is still worse on GEC benchmarks than fine-tuned sequence-to-sequence models, both in the zero-shot and few-shot scenarios. [Coyne et al. \(2023\)](#) provide the corresponding analysis for GPT-4, with even lower results, which means that specialized models are still relevant for GEC scenarios and validate our research.

3 Approach

In this section, we introduce our approach that uses multi-task learning and optimizes the training schedule for the sequence-to-sequence model architecture. In Section 3.1, we outline the model and hyperparameters used for training. In Section 3.2, we describe existing GEC datasets for training and evaluation, highlighting the specific properties of each. In Section 3.3, we specify the main and auxiliary tasks formulated as sequence-to-sequence mapping. Section 3.4 describes the training steps.

3.1 Model

We use a straightforward text-only approach while also keeping the model size limited. Namely, we formulate all training tasks as sequence-to-sequence text rewriting problems and do not introduce any additional heads for edit prediction or other tasks. As the backbone, we use the BART ([Lewis et al., 2020](#)) model with a 12-layer encoder and decoder. We train the model in fp16 mode with learning rate $1e^{-5}$ and warmup ratio 7% with a linear scheduler. We fit 12 instances in a batch and use gradient accumulation with 4 steps to achieve a larger effective batch size. We did not perform a hyperparameter search except for the learning rate, which was chosen from $[1e^{-5}, 5e^{-5}, 1e^{-6}]$. All training experiments were done on 8 NVIDIA Tesla V100 GPUs with 32GB of memory.

3.2 Training Data

For pretraining (Stage I), we use C4_{200M} or PIE datasets. C4_{200M} is a synthetic corpus based on clean sentences from the C4 dataset. This corpus was generated by a tagged corruption model to

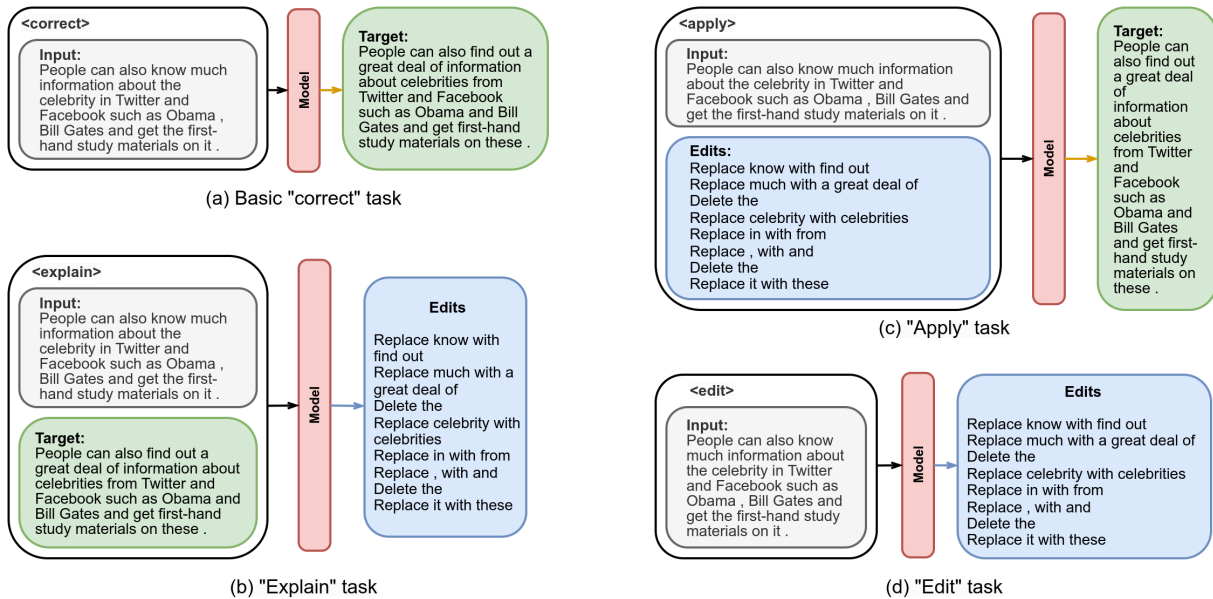


Figure 1: Tasks automatically generated from a source-target pair for multi-task pretraining.

meet the error distribution of BEA-dev (Bryant et al., 2019a); see Stahlberg and Kumar (2021) for details. PIE is a synthetic dataset of 9M parallel sentences generated by using rule-based grammatical errors such as deletion, insertion, and word replacement (Awasthi et al., 2019a).

For other training stages, we use the following datasets: (i) National University of Singapore Corpus of Learner English (NUCLE) (Dahlmeier et al., 2013) that consists of essays written by undergraduate students on different topics and annotated by professional English instructors; (ii) Lang-8 Corpus of Learner English (Lang-8) (Tajiri et al., 2012) collected from the online language learning site Lang-8; this dataset is relatively “noisy” as the users corrected themselves, and it comes with several annotations; (iii) First Certificate in English (FCE) (Yannakoudakis et al., 2011) with short texts written by English learners as answers to exam questions assessing the upper-intermediate level; this dataset is relatively clean but covers only a single group of English learners; (iv) Write & Improve + LOCNESS Corpus (W&I+L) (Bryant et al., 2019a); Write & Improve dataset consists of text blocks (essays, letters etc.) written by English learners and submitted to the W&I system; LOCNESS is composed of essays written by native English students, used only for evaluation; these datasets are the “cleanest” and have a wide distribution over different levels of English. Here and below, *errorful* sentences are those that contain at least one error. We use the BEA-2019 devel-

opment set, i.e. W&I+L-dev, to choose the best model and report results on the CoNLL2014 test set (Ng et al., 2014) evaluated by the official M2 scorer (Dahlmeier and Ng, 2012) and the BEA-2019 test set evaluated by ERRANT (Bryant et al., 2017). Table 1 summarizes dataset statistics and shows which training stages they are used for.

3.3 Multi-task learning

The standard approach is to view GEC as a sequence-to-sequence mapping from errorful sentences to corrected ones. Since the sentences are from the same language, one can align two sentences using an edit distance algorithm. The alignment can be transformed into a sequence of *insert*, *delete*, and *replace* operations that transform the errorful sentence into a correct one. We find the list of operations using ERRANT² (Bryant et al., 2019b) and use them for auxiliary tasks.

Inspired by recent works on chain of thought prompting (Wei et al., 2023; Zhou et al., 2022), we construct several tasks and examine their influence on the model’s performance. Each task, as illustrated in Fig. 1, is explicitly marked by a prefix written as $\langle prefix \rangle$ and followed by an input string:

- (i) **Correct:** standard generation of a corrected sentence given the original, encoded as “ $\langle correct \rangle source$ ” and trained to produce the target sentence (Fig. 1a);

²We use version 2.3 of ERRANT for both evaluation and training set construction.

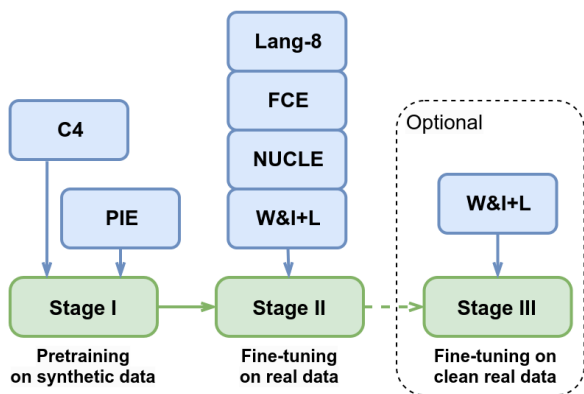


Figure 2: Training order for a GEC model.

- (ii) **Explain:** generation of a sequence of explanations (atomic edits) given both original and corrected sentences; encoded as “*<explain>* Input: *src \n Target: tgt*” with the result in the format “*Delete smth \n Insert smth...*” (Fig. 1b); if no correction is needed, the model should generate “*No correction*”;
- (iii) **Apply:** application of edits to the original sentence to get the target, encoded as “*<apply>* Input: *src \n Do: Delete smth \n Insert smth \n Replace smth with smth*”; the result should be the correct target sentence (Fig. 1c);
- (iv) **Edit:** generation of the sequence of edits given the original sentence but not the target, encoded as “*<edit> src*” with the target in the form “*Delete smth \n Insert smth \n Replace smth with smth*” (Fig. 1d); if no correction is needed, the model should generate “*No correction*”.

We generate an auxiliary sequence-to-sequence pair for every task and for every pair from the original dataset.

3.4 Training Order

Tarnavskiy et al. (2022) and Omelanchuk et al. (2021) mention that training order is essential to obtain the best possible results. They separate the overall process into three stages: pretrain the model on synthetic data, fine-tune it on errorful sentences from the four GEC datasets: LANG-8, NUCLE, FCE, and W&I+L, and then fine-tune it on the clean GEC dataset W&I+L.

In this work, we suggest to modify this procedure and claim that the training process benefits from choosing a correct ordering of data. Our process is illustrated in Fig. 2: on Stage I, we similarly

pretrain our model on large-scale synthetic data. We consider two datasets that differ both in size and in generation approach: $C4_{200M}$ and PIE. This step adapts the model for the downstream task. On Stage II, we fine-tune on four GEC datasets but modify the above procedure. First, we use all sentences, not only errorful ones. Second, we use the datasets in a strict order: Lang-8, NUCLE, FCE, W&I+L, with no shuffling across datasets. Third, we do not shuffle samples inside the datasets either: each dataset is composed of coherent texts, so we thus preserve their original structure and place examples from the same texts together.

For a more complete comparison with previous works, we add Stage III where we additionally fine-tune the model on the W&I+L dataset. We note that this step helps to increase recall without a substantial decrease in precision, yielding modest improvements in $F_{0.5}$. Note that in previous works, this step was mandatory as the target distribution is correlated with the W&I+L dataset. In contrast, we add this dataset as the last in our training order, which is a more natural approach; the suggested scheme also looks more suitable for real world tasks where there is no obvious way to split the data into different stages.

3.5 Re-weighted Sampling

The $F_{0.5}$ metric is commonly used for the evaluation of GEC models. It puts a higher weight on the model’s precision than recall. In our experiments, we see that the proposed training pipeline makes the model less aggressive in editing which improves performance. Another approach to making the model choose corrections that it is confident about is the Align-Pred method proposed by Sun and Wang (2022). It increases the probability of the next original token (obtained by aligning the original and currently generated sequences) during decoding. In order to show that our method is orthogonal to this, we apply Align-Pred to our best model. We introduce a modification to this method that goes beyond the original, applying temperature scaling before Align-Pred. This significantly improves Align-Pred (see Table 2).

4 Evaluation and Analysis

4.1 Comparison with baselines

We compare the proposed model with several state of the art baselines. We used three variations of the tagging model GECToR: with the XLNet-L back-

Model	Size	Synthetic	CoNLL-14			BEA-test		
			Prec	Rec	F _{0.5}	Prec	Rec	F _{0.5}
T5 (Rothe et al., 2021)	770M	mC4(>300M)+CLANG	—	—	66.04	—	—	72.06
T5-XL (Rothe et al., 2021)	3B	mC4(>300M)+CLANG	—	—	67.65*	—	—	73.92
T5-XXL (Rothe et al., 2021)	11B	mC4(>300M)+CLANG	—	—	68.75*	—	—	<u>75.88</u>
GECToR-XLNet (Omelianchuk et al., 2020)	350M	PIE(9M)	77.5	40.1	65.3	79.2	53.9	72.4
GECToR-RoBERTa (Tarnavskiy et al., 2022)	350M	PIE(9M)	74.40	41.05	64.0	80.70	53.39	73.21
GECToR-XLNet (Lai et al., 2022)	350M	PIE(9M)	78.18	42.67	67.02	81.33	51.55	72.91
Transformer (Stahlberg and Kumar, 2021)	340M	C ₄ _{200M}	72.8	<u>49.5</u>	66.6	72.1	64.4	70.4
BART-based (Sun and Wang, 2022)	400M	300M	—	—	—	76.1	65.6	73.8
BART-based + Align-and-Predict (Sun and Wang, 2022)	400M	300M	—	—	—	78.7	63.2	75.0
BART (ours)	400M	C ₄ _{200M}	75.43	51.20	<u>68.91</u>	78.19	<u>65.54</u>	75.28
BART (ours) + Align-and-Predict + Temp. Scaling	400M	PIE(9M)	75.45	47.03	67.31	78.49	<u>58.66</u>	73.52
BART (ours) + Align-and-Predict + Temp. Scaling	400M	C ₄ _{200M}	<u>78.00</u>	49.12	69.79	<u>80.88</u>	61.15	75.97

Table 2: Evaluation results on CoNLL-14 and BEA-test. Best results are shown in **bold**, second best are underlined; * results shown by the *arXiv* version (Rothe et al., 2022) that differs from (Rothe et al., 2021) on the CONLL test set

Pre-training	Stage II	Stage III	Prec	Rec	F _{0.5}
C ₄ _{200M}	—	—	36.67	24.83	33.88
C ₄ _{200M}	4 datasets	W&I+L	66.77	50.14	62.62
PIE	—	—	44.55	20.04	35.79
PIE	4 datasets	W&I+L	65.00	46.29	60.14

Table 3: Influence of the pretraining dataset. Evaluation results on BEA-dev.

bone (Omelianchuk et al., 2020), with RoBERTa backbone and multi-stage training (Tarnavskiy et al., 2022), and with the XLNet-L backbone and multiturn corrections (Lai et al., 2022). For *seq2seq* baselines we consider the vanilla GEC Transformer trained on C₄_{200M} (Stahlberg and Kumar, 2021), T5-XL and T5-XXL models (Rothe et al., 2021) also trained on the C4 dataset, and a BART-based model trained on 300M synthetic data (Sun and Wang, 2022). All baselines have comparable or larger model sizes than ours and use the same PIE synthetic dataset or larger synthetic collections.

Table 2 shows evaluation results for the CoNLL-14 (Ng et al., 2014) and BEA-test (Bryant et al., 2019a) datasets. Our approach shows the best results for comparable model sizes with a significant margin and outperforms all models, including T5-XL with 3B parameters and even T5-XXL which is 30x larger than our model (11B parameters) and trains on a larger synthetic dataset. We also present evaluation results of our approach with pretraining on the PIE dataset to compare with the models pretrained on it. Again, we see that our model outperforms more sophisticated methods such as (Lai et al., 2022).

4.2 Influence of the Pretraining Dataset

In this section, we analyze the choice of the pretraining dataset, comparing two publicly available synthetic datasets: C₄_{200M} and PIE. They differ not only in size (see Table 1) but also in the type of generated errors that are model-based and rule-based respectively. Table 3 shows the performance of models pretrained in different setups. The model pretrained on C₄_{200M} has higher recall indicating that it covers more errors, while the model pretrained on PIE reaches higher precision. Note that almost all sentences from synthetic datasets contain errors, which means that the pretraining distribution differs a lot from the development set. Hence, to make a fair comparison we further fine-tune the models on GEC-specific datasets using the standard three-stage approach. Table 3 shows that the model pretrained on C₄_{200M} performs better in terms of precision, recall, and the F_{0.5} score.

4.3 Order of the Datasets

We also examine the influence of the ordering of training datasets and examples within a dataset. We are mainly concerned with multi-task training, but note that the training schedule also impacts the single-task pipeline (see Table 4). The model trained using a specific schedule achieves good results after the first stage of fine-tuning, while the third stage improves the model’s recall and makes its results surpass the standard three-stage training.

For the multi-task approach, we use three tasks—Correct, Explain, and Apply—and different dataset schedules. There are 16 possible orderings, but Lang-8 is a “noisy” dataset so it is reasonable to place it early in the fine-tuning, while W&I+L is

Training approach	CoNLL-14			BEA-test			BEA-dev		
	Prec	Rec	F _{0.5}	Prec	Rec	F _{0.5}	Prec	Rec	F _{0.5}
Two stages	66.47	47.65	61.60	59.61	63.74	60.39	47.94	40.91	46.35
Three stages	73.64	<u>53.16</u>	68.37	75.65	68.94	74.20	66.77	<u>50.14</u>	<u>62.62</u>
Two stages, optimized schedule	77.18	48.52	69.02	78.71	63.05	<u>74.90</u>	<u>68.60</u>	41.56	60.70
Three stages, optimized schedule	74.96	52.34	<u>69.00</u>	77.17	66.39	74.74	68.15	46.31	62.28
Multi-task: three tasks, two stages	66.51	48.47	61.79	59.55	63.53	60.31	48.22	41.50	46.71
Multi-task: three tasks, three stages	73.96	53.30	68.64	75.85	<u>68.63</u>	74.29	65.62	52.07	62.38
Multi-task: three tasks, optimized schedule	<u>75.43</u>	51.20	68.91	<u>78.19</u>	65.54	75.28	68.64	46.50	62.67

Table 4: Comparison of standard and multi-task fine-tuning; all models are pretrained on C4_{200M}. Best results are shown in **bold**, second best are underlined.

Order	Prec	Rec	F _{0.5}
Order preserved within each dataset			
FCE → Lang-8 → NUCLE → W&I+L	68.72	44.98	62.15
NUCLE → Lang-8 → FCE → W&I+L	68.86	45.02	62.26
Lang-8 → FCE → NUCLE → W&I+L	68.52	45.32	62.16
Lang-8 → NUCLE → FCE → W&I+L	68.78	45.98	62.67
FCE → Lang-8 → W&I+L → NUCLE	77.41	13.70	40.11
Instances shuffled within each dataset			
FCE → Lang-8 → NUCLE → W&I+L	66.41	48.47	61.84
Lang-8 → NUCLE → FCE → W&I+L	66.99	48.48	62.24
Instances shuffled across datasets			
Lang-8, NUCLE, FCE, W&I+L	48.21	41.20	46.63

Table 5: Ablation study of the dataset fine-tuning order. The model is trained in 2-stages, evaluated on BEA-dev.

Stage II	Stage III	Prec	Rec	F _{0.5}
Order preserved within each dataset				
FCE → Lang-8 → NUCLE → W&I+L	W&I+L	67.88	47.02	62.34
NUCLE → Lang-8 → FCE → W&I+L	W&I+L	68.23	47.12	62.62
Lang-8 → FCE → NUCLE → W&I+L	W&I+L	68.14	47.61	62.73
Lang-8 → NUCLE → FCE → W&I+L	W&I+L	67.35	49.89	62.94
FCE → Lang-8 → W&I+L → NUCLE	W&I+L	67.66	47.32	62.30
Instances shuffled within each dataset				
FCE → Lang-8 → NUCLE → W&I+L	W&I+L	66.05	50.03	62.08
Lang-8 → NUCLE → FCE → W&I+L	W&I+L	66.05	50.34	62.17
Instances shuffled across datasets				
Lang-8, NUCLE, FCE, W&I+L	W&I+L	65.62	52.07	62.38

Table 6: Ablation study of the dataset training order. The model is trained in 3-stages. Evaluation is done on the BEA-dev dataset.

“cleaner” than others so we use it late in the training. Tables 5 and 6 shows the performance of models trained in different setups. It shows that, first, the order of datasets really matters: some cases outperform others by a significant margin. Second, shuffling instances within a dataset also reduces performance, which supports our hypothesis that sentences from the same block of text should be placed together while training. Third, the best

Tasks	Prec	Rec	F _{0.5}
Correct + Explain	69.03	44.20	62.06
Correct + Apply	68.64	46.50	62.67
Correct + Edit	69.19	44.92	62.45
Correct + Explain + Apply	68.78	45.98	62.57
Correct + Apply + Edit	68.98	44.95	62.32
Correct + Explain + Apply + Edit	68.58	45.81	62.38
Edit + Apply	49.41	28.80	43.22

Table 7: Ablation study on auxiliary tasks. Evaluation is done on the BEA-dev dataset.

performance is achieved when W&I+L is the last dataset in the schedule, as evidenced by the drop in performance for the “FCE → Lang-8 → W&I+L → NUCLE” setup and improvement after further fine-tuning on the W&I+L dataset.

4.4 Multi-Task Pretraining

Table 4 compares standard fine-tuning with the proposed multi-task fine-tuning. We show the commonly used metrics—precision, recall, and F_{0.5} score—on three GEC datasets: CoNLL-14, BEA-test, and BEA-dev. In all cases, multi-task training leads to improved performance for the final task, especially for the BEA-dev dataset.

Table 7 presents an ablation study related to the choice of auxiliary tasks; we fix the training schedule and use different task compositions. We make the following observations. First, the best combination of tasks is “Correct” and “Apply”, yielding the highest recall and high enough precision compared to other settings. Second, adding the edit prediction task (“Edit”) in combination with other tasks lowers the model’s performance; below, we argue that this could be due to the complexity of the task that the model cannot effectively learn. On the other hand, while the “Edit” task does not improve performance, it could help to make the model interpretable without using external tools.

Additionally, we study the case where we replace

the GEC task with consecutive “Edit” and “Apply” tasks, decomposing GEC into a chain-of-thought process: first the model says what is wrong and what actions should be taken and then it applies these actions to the original sentence to produce the final output. Table 7 shows that the model trained with only these two tasks severely underperforms, so “classical” GEC is still necessary.

4.5 Auxiliary Tasks Analysis

Next, we examine the model’s performance on auxiliary tasks. For the analysis of “Explain” and “Apply” tasks, we use the model that had them both in pretraining; to analyze the “Edit” task, we use a model pretrained on all four tasks. We use the CoNLL-14 dataset for evaluation.

The “Explain” task urges the model to align to sentences and produce a correct sequence of edits to transform the source into the target. The exact match score of generated edits is 90% compared to the gold standard edits, and the $F_{0.5}$ score is 91.55. Moreover, many “errors” here come from the ambiguity of how to perform replacements or from merging/splitting corrections compared to the gold standard. For example, for the original sentence “*People can also know much information about the celebrity in Twitter and Facebook such as Obama , Bill Gates and get the first-hand study materials on it .*” an annotator suggested the following corrections: replace *know much* with *find out a great deal of*; delete *the*; replace *celebrity* with *celebrities*; replace *in* with *from*; replace *,* with *and*; delete *the*; replace *it* with *these*. The model correctly predicts all of them except one. It splits the first correction into two: replace *know* with *find out* and replace *much* with *a great deal*.

The model deals quite well with the “Apply” task, getting 93.5% accuracy. Most errors correspond to doing a different operation (e.g., inserting a new word instead of replacing one word with another), applying edits in the wrong place (e.g., placing new tokens after a wrong word), and simply ignoring some of the edits, especially if the list of edits is long. It appears that some errors arise from ambiguous descriptions of edits—in our approach we specify edits as, e.g., “Insert the”, and the model has to decide where to insert it by itself—so performance could be improved further. For an example, consider the following sentence: *It is a concern that will be with us during our whole life , because we will never know when the ”potential*

bomb’ ’ will explode . Here, the correction “delete *will*” targets the second *will*, but the model does not make changes to the original sentence, ignoring the correction. In this work, we restricted ourselves to simplistic prompts, and we leave an exploration of more detailed edit descriptions for further work.

The last auxiliary task, “Edit”, is the prediction of a list of edits. This task is hard for the model: the exact match score is only 23.5%, and the $F_{0.5}$ score is 30.69. This low performance might be the reason why adding this task to the training pipeline decreases GEC performance. There are many cases where the model correctly applies an edit using a prediction prompt but omits it with the “Edit” prompt, which could indicate that the two tasks do not interact properly. It seems that the model struggles to predict a combination of connected or related corrections. For example, for the sentence “*The people with albinism have sensitive skin and it needs regular treatment .*”, annotators gave the following edits: replace *The people* with *People*; replace *it* with *this*. But the model predicts the following: delete *The*; replace *it* with *they*. Thus, after correctly removing *the* it does not capitalize *people*, and while replacing *it* with *they* it does not replace *needs* with *need*.

Since our model can be trained on both “Edit” and “Apply” tasks, it is tempting to apply a chain-of-thought procedure where we first predict the list of operations and then apply them to the original sentence. Applying such a procedure, we get the $F_{0.5}$ score of 52.18 on the CoNLL-14 dataset (precision 55.85, recall 41.33), i.e., the most problematic metric is precision (see also Table 7). The main problem of this task is the ambiguity of the target sentence and edits needed to obtain it: edits can interact with each other, which is hard to account for given only the source sentence. This chain-of-thought approach appears unnecessarily challenging, while adding the target sentence in the “Explain” task makes the problem much easier and more helpful for pretraining.

We have also studied how tasks interact with each other. In the first experiment, we correct a sentence with the model and then ask it to explain the corrections. We compare the explanations with edits produced by the ERRANT tool, obtaining an exact match score of 95.9%, higher by 5.9% than for human-corrected sentences. This indicates some interaction between tasks as the model’s corrections are more explainable by itself. Second, we

chain three tasks: first the model corrects a sentence, then it explains the corrections, and finally it applies edits to the original sentence. Interestingly, the exact match between the corrected sentence on the first step and the sentence obtained after chaining is 97.02%, so the model is not always consistent across the tasks. Many errors come from the ambiguity of edits. It may be possible to use the discrepancy between corrected and generated sentences for filtering, either leaving the initial sentence intact or taking the intersection of the edits related to the two sentences. We find, however, that these approaches do not improve the final performance or even decrease it.

4.6 Automatic Human Evaluation

In this part, we compare the performance of our model with ChatGPT with chain-of-thoughts prompting, using the results by Fang et al. (2023) who showed that ChatGPT performs poorly on common benchmarks, far below state of the art fine-tuned models in the $F_{0.5}$ metric, with high recall but very low precision. They show that the model often proposes corrections that are reasonable but far from minimal with respect to the number of required edits. To punish such over-corrections less, Bryant and Ng (2015) propose to use a test set where each erroneous sentence is annotated by several experts and averages the scores obtained in comparison to each annotator; this approach increases the scores for edits proposed by at least one annotator. Following Fang et al. (2023), we use the evaluation set from this paper consisting of CoNLL-14 sentences with 10 human annotations each. We also compare with a human-level baseline which is measured as the average score of each human annotator with respect to others. Table 8 shows that even in this setup our model outperforms ChatGPT by a large margin; we compare with Fang et al. (2023) who use chain-of-thought prompting and report results in the zero-shot and few-shot (1, 3, and 5) settings. Interestingly, ChatGPT performance is slightly below the human level, while our model performs significantly better.

5 Conclusion

In this work, we have introduced a new approach to training and fine-tuning sequence-to-sequence models for the grammatical error correction task based on multi-task pretraining and optimized training schedules. The proposed approach fares better

System	$F_{0.5}$
Human level	72.58
Transformer	66.97
GECToR	80.49
T5-large	81.19
ChatGPT (0-shot)	69.74
ChatGPT (1-shot)	71.55
ChatGPT (3-shot)	71.73
ChatGPT (5-shot)	70.66
Our	81.36

Table 8: Comparison to automatic human evaluation performance on 10 references for CoNLL-14; ChatGPT results are taken from Fang et al. (2023).

than previous state of the art models with a much smaller model size and outperforms models of comparable size by a wide margin on both CoNLL-14 and BEA-test datasets; in particular, we have achieved results exceeding state of the art models based on T5-XXL (11B parameters) with a BART-based model (400M parameters). Our multi-task approach encodes auxiliary tasks as text rewriting problems and does not require any changes in the model architecture. We believe that the proposed approach is of significant value for practical GEC pipelines and see several directions for possible improvements in further work, including modifications of text encodings for auxiliary tasks and adding new auxiliary tasks to the training process.

6 Limitations

Limitations of our study can serve as motivation for further work. First, we do not test our approach in the increasingly important multilingual setup; we believe that our ideas may contribute even more to multilingual models that need greater robustness. Second, the highest possible results in GEC are currently obtained by ensembles rather than individual models, including the recent state of the art approach by Qorib et al. (2022); while ensembling would make our reduced model size less important, it would be interesting to study how well our model can perform in such ensembles. Third, we consider only simple prompts and small models, and our results could be extended further.

Acknowledgements

The work of Sergey Nikolenko was done under support of the grant no. 075-15-2022-289 for the creation and development of Euler International Mathematical Institute.

References

- Abhijeet Awasthi, Sunita Sarawagi, Rasna Goyal, Sabyasachi Ghosh, and Vihari Piratla. 2019a. [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.
- Abhijeet Awasthi, Sunita Sarawagi, Rasna Goyal, Sabyasachi Ghosh, and Vihari Piratla. 2019b. [Parallel iterative edit models for local sequence transduction](#). *ArXiv*, abs/1910.02893.
- Chris Brockett, William B. Dolan, and Michael Gamon. 2006. [Correcting ESL errors using phrasal SMT techniques](#). In *Proceedings of the 21st International Conference on Computational Linguistics and 44th Annual Meeting of the Association for Computational Linguistics*, pages 249–256, Sydney, Australia. Association for Computational Linguistics.
- Christopher Bryant, Mariano Felice, Øistein E. Andersen, and Ted Briscoe. 2019a. [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, Florence, Italy. Association for Computational Linguistics.
- Christopher Bryant, Mariano Felice, Øistein E. Andersen, and Ted Briscoe. 2019b. [The bea-2019 shared task on grammatical error correction](#). In *BEA@ACL*.
- Christopher Bryant, Mariano Felice, and Ted Briscoe. 2017. [Automatic annotation and evaluation of error types for grammatical error correction](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 793–805, Vancouver, Canada. Association for Computational Linguistics.
- Christopher Bryant and Hwee Tou Ng. 2015. [How far are we from fully automatic high quality grammatical error correction?](#) 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 697–707, Beijing, China. Association for Computational Linguistics.
- Yo Joong Choe, Jiyeon Ham, Kyubyong Park, and Yeol Yoon. 2019. [A neural grammatical error correction system built on better pre-training and sequential transfer learning](#). In *Proceedings of the Fourteenth Workshop on Innovative Use of NLP for Building Educational Applications*, pages 213–227, Florence, Italy. Association for Computational Linguistics.
- Steven Coyne, Keisuke Sakaguchi, Diana Galvan-Sosa, Michael Zock, and Kentaro Inui. 2023. [Analyzing the performance of gpt-3.5 and gpt-4 in grammatical error correction](#).
- 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, Montréal, Canada. Association for Computational Linguistics.
- 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, Atlanta, Georgia. Association for Computational Linguistics.
- Meiyuan Fang, Kai Fu, Jiping Wang, Yang Liu, Jin Huang, and Yitao Duan. 2020. [A hybrid system for NLPTEA-2020 CGED shared task](#). In *Proceedings of the 6th Workshop on Natural Language Processing Techniques for Educational Applications*, pages 67–77, Suzhou, China. Association for Computational Linguistics.
- Tao Fang, Shu Yang, Kaixin Lan, Derek F. Wong, Jinpeng Hu, Lidia S. Chao, and Yue Zhang. 2023. [Is chatgpt a highly fluent grammatical error correction system? a comprehensive evaluation](#).
- Jennifer Foster and Oistein Andersen. 2009. [GenERRate: Generating errors for use in grammatical error detection](#). In *Proceedings of the Fourth Workshop on Innovative Use of NLP for Building Educational Applications*, pages 82–90, Boulder, Colorado. Association for Computational Linguistics.
- 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, Florence, Italy. Association for Computational Linguistics.
- Jiatao Gu, Changhan Wang, and Jake Zhao. 2019. [Levenshtein transformer](#). In *Neural Information Processing Systems*.
- Phu Mon Htut and Joel Tetreault. 2019. [The unbearable weight of generating artificial errors for grammatical error correction](#). In *Proceedings of the Fourteenth Workshop on Innovative Use of NLP for Building Educational Applications*, pages 478–483, Florence, Italy. Association for Computational Linguistics.
- Masahiro Kaneko, Kengo Hotate, Satoru Katsumata, and Mamoru Komachi. 2019. [TMU transformer system using BERT for re-ranking at BEA 2019 grammatical error correction on restricted track](#). In *Proceedings of the Fourteenth Workshop on Innovative Use of NLP for Building Educational Applications*, pages 207–212, Florence, Italy. Association for Computational Linguistics.

- Satoru Katsumata and Mamoru Komachi. 2020. [Stronger baselines for grammatical error correction using a pretrained encoder-decoder model](#). In *Proceedings of the 1st Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 10th International Joint Conference on Natural Language Processing*, pages 827–832, Suzhou, China. 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.
- Shaopeng Lai, Qingyu Zhou, Jiali Zeng, Zhongli Li, Chao Li, Yunbo Cao, and Jinsong Su. 2022. [Type-driven multi-turn corrections for grammatical error correction](#). In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 3225–3236, Dublin, Ireland. Association for Computational Linguistics.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. [BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880, Online. Association for Computational Linguistics.
- Jared Lichtarge, Chris Alberti, and Shankar Kumar. 2020. [Data weighted training strategies for grammatical error correction](#). *Transactions of the Association for Computational Linguistics*, 8:634–646.
- Masato Mita, Shun Kiyono, Masahiro Kaneko, Jun Suzuki, and Kentaro Inui. 2020. [A self-refinement strategy for noise reduction in grammatical error correction](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 267–280, Online. Association for Computational Linguistics.
- Jakub Náplava and Milan Straka. 2019. [Grammatical error correction in low-resource scenarios](#). In *Proceedings of the 5th Workshop on Noisy User-generated Text (W-NUT 2019)*, pages 346–356, Hong Kong, China. Association for Computational Linguistics.
- 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, Baltimore, Maryland. Association for Computational Linguistics.
- 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, Seattle, WA, USA → Online. Association for Computational Linguistics.
- Kostiantyn Omelianchuk, Vipul Raheja, and Oleksandr Skurzhanskyi. 2021. [Text Simplification by Tagging](#). In *Proceedings of the 16th Workshop on Innovative Use of NLP for Building Educational Applications*, pages 11–25, Online. Association for Computational Linguistics.
- Muhammad Qorib, Seung-Hoon Na, and Hwee Tou Ng. 2022. [Frustratingly easy system combination for grammatical error correction](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1964–1974, Seattle, United States. Association for Computational Linguistics.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, Peter J Liu, et al. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *J. Mach. Learn. Res.*, 21(140):1–67.
- Sascha Rothe, Jonathan Mallinson, Eric Malmi, Sebastian Krause, and Aliaksei Severyn. 2021. [A simple recipe for multilingual grammatical error correction](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 702–707, Online. Association for Computational Linguistics.
- Sascha Rothe, Jonathan Mallinson, Eric Malmi, Sebastian Krause, and Aliaksei Severyn. 2022. [A simple recipe for multilingual grammatical error correction](#).
- Felix Stahlberg and Shankar Kumar. 2021. [Synthetic data generation for grammatical error correction with tagged corruption models](#). In *Proceedings of the 16th Workshop on Innovative Use of NLP for Building Educational Applications*, pages 37–47, Online. Association for Computational Linguistics.
- Xin Sun and Houfeng Wang. 2022. [Adjusting the precision-recall trade-off with align-and-predict decoding for grammatical error correction](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 686–693, Dublin, Ireland. Association for Computational Linguistics.
- Toshikazu Tajiri, Mamoru Komachi, and Yuji Matsumoto. 2012. [Tense and aspect error correction for ESL learners using global context](#). In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 198–202, Jeju Island, Korea. Association for Computational Linguistics.

- Maksym Tarnavskiy, Artem Chernodub, and Kostiantyn Omelianchuk. 2022. [Ensembling and knowledge distilling of large sequence taggers for grammatical error correction](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3842–3852, Dublin, Ireland. Association for Computational Linguistics.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed Chi, Quoc Le, and Denny Zhou. 2023. [Chain-of-thought prompting elicits reasoning in large language models](#).
- Konstantin Yakovlev, Alexander Podolskiy, Andrey Bout, Sergey Nikolenko, and Irina Piontkovskaya. 2023. [GEC-DePenD: Non-autoregressive grammatical error correction with decoupled permutation and decoding](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1546–1558, Toronto, Canada. Association for Computational Linguistics.
- 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, Portland, Oregon, USA. Association for Computational Linguistics.
- Zheng Yuan, Felix Stahlberg, Marek Rei, Bill Byrne, and Helen Yannakoudakis. 2019. [Neural and FST-based approaches to grammatical error correction](#). In *Proceedings of the Fourteenth Workshop on Innovative Use of NLP for Building Educational Applications*, pages 228–239, Florence, Italy. Association for Computational Linguistics.
- Zheng Yuan, Shiva Taslimipoor, Christopher Davis, and Christopher Bryant. 2021. [Multi-class grammatical error detection for correction: A tale of two systems](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 8722–8736, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Yi Zhang, Tao Ge, Furu Wei, Ming Zhou, and Xu Sun. 2019. Sequence-to-sequence pre-training with data augmentation for sentence rewriting. *arXiv preprint arXiv:1909.06002*.
- Yue Zhang, Bo Zhang, Zhenghua Li, Zuyi Bao, Chen Li, and Min Zhang. 2022. [SynGEC: Syntax-enhanced grammatical error correction with a tailored GEC-oriented parser](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 2518–2531, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Denny Zhou, Nathanael Schärli, Le Hou, Jason Wei, Nathan Scales, Xuezhi Wang, Dale Schuurmans, Olivier Bousquet, Quoc Le, and Ed Chi. 2022. Least-to-most prompting enables complex reasoning in large language models. *arXiv preprint arXiv:2205.10625*.

A Error Type Distribution in Datasets

In order to analyze the difference between the GEC datasets, we compute the error type statistics using the ERRANT tool. Next, we compare each dataset with W&I+L-dev, see Figures 4-9. We see that W&I+L-train and W&I+L-dev have only a slight discrepancy in the proportions of each error type, with major differences coming from punctuation and spelling errors. Hence, we should expect that the distribution on W&I+L-test is close but not exactly the same as the articles that comprise the datasets were written by different people. Thus, extensive hyperparameter search on W&I+L-dev may lead to lower performance. The FCE dataset is also close to these datasets except for SPELL, NOUN, and VERB errors. These differences can be explained by the fact that essays were written by authors who had only begun learning English, and these error types are more common for them.

Comparing W&I+L-dev with NUCLE and Lang-8 datasets, we note that for some error types the difference is striking. We highlight that the proportion of OTHER errors is high for both datasets. In the ERRANT toolkit, OTHER corresponds to errors that cannot be labeled by a predefined set of error types. After examining some instances with that error type, we have noticed that a large part of them is related to sentence rephrasing or rewriting, perhaps regrouping some parts of the sentence. Training on such examples might hurt the model’s performance because it would become more aggressive in performing corrections, leading to much lower precision. Therefore, we mark those datasets as “noisy” and others as “clean”.

The combination of all datasets—FCE, W&I+L-train, Lang-8, and NUCLE—does not improve the situation, as shown in Figure 9. Nevertheless, if we train only on FCE and W&I+L-train, we obtain poor performance. This indicates that not only the distribution of errors should be close but also the errors should be diverse in order to generalize well. These two factors reveal why we needed multi-stage fine-tuning.

Our approach can be viewed as curriculum learning where we gradually train a model on “less noisy” data.

B Model comparison

Performance on the W&I+L-dev dataset of the 3-stage model and of the multi-task model with the improved training schedule seems to be marginal:

62.62 versus 62.67. However, the models’ behavior on W&I+L-dev differs. The 3-stage model has higher recall and the other has higher precision. As we have noted in the previous section, error distributions on parts of W&I+L datasets differ from each other since different people are prone to making different types of errors. Therefore, we expect that a model with higher precision that makes more accurate corrections would generalize better. Looking at the models’ performance on W&I+L-test and CoNLL-14, we see that the gap between the models is indeed large.

To draw the distinction further, we compute the statistics of error types corrected and introduced by them with the ERRANT toolkit. Figure 10 presents the absolute number of corrected, generated, and non-corrected errors for every type. Again, we see that the 3-stage model is more aggressive. It corrects more errors of each type but also introduces more new errors. In Figures 11-13, we show the distribution of corrected, induced, and non-corrected errors by type. Again, we see that the models’ behavior differs.




Figure 3: W&I+L-dev and W&I+L-train error type comparison.




Figure 4: W&I+L-dev and NUCLE error type comparison.




Figure 5: W&I+L-dev and LANG-8 error type comparison.




Figure 6: W&I+L-dev and FCE error type comparison.




Figure 7: W&I+L-dev and CoNLL-14 error type comparison.




Figure 8: NUCLE and CoNLL-14 error type comparison.




Figure 9: W&I+L-dev and BEA-train error type comparison.




Figure 10: 3-stage and multi-task improved schedule models comparison. Corrected — errors corrected by the model. Generated — errors introduced by the model. Left — errors that were not corrected.




Figure 11: 3-stage and multi-task improved schedule models corrected errors distribution by type.




Figure 12: 3-stage and multi-task improved schedule models introduced errors distribution by type.




Figure 13: 3-stage and multi-task improved schedule models non-corrected errors distribution by type.