

Small Language Models Improve Giants by Rewriting Their Outputs

Giorgos Vernikos^{1,2*} Arthur Bražinskas³ Jakub Adamek³
Jonathan Mallinson³ Aliaksei Severyn³ Eric Malmi³

¹EPFL, ²HEIG-VD / HES-SO, ³Google Research

georgios.vernikos@epfl.ch

{abrazinskas, enkait, jonmall, severyn, emalmi}@google.com

Abstract

Despite the impressive performance of large language models (LLMs), they often lag behind specialized models in various tasks. LLMs only use a fraction of the existing training data for in-context learning, while task-specific models harness the full dataset for fine-tuning. In this work, we tackle the problem of leveraging training data to improve the performance of LLMs without fine-tuning. Our approach directly targets LLM predictions without requiring access to their weights. We create a pool of candidates from the LLM through few-shot prompting and we employ a compact model, the LM-corrector (LMCOR), specifically trained to merge these candidates to produce an enhanced output. Our experiments on four natural language generation tasks demonstrate that even a small LMCOR model (250M) substantially improves the few-shot performance of LLMs (62B), matching and even outperforming standard fine-tuning. Furthermore, we illustrate the robustness of LMCOR against different prompts, thereby minimizing the need for extensive prompt engineering. Finally, we show that LMCOR can be seamlessly integrated with different LLMs at inference, serving as a plug-and-play module to improve their performance.

1 Introduction

Large language models have recently demonstrated near state-of-the-art performance on various tasks via in-context learning, which enables them to generate outputs based on instructions and a handful of examples, without task-specific training (Brown et al., 2020b,a; Chowdhery et al., 2022). However, the effectiveness of this paradigm can vary significantly depending on the task instruction (Shin et al., 2020; Jiang et al., 2021; Schick and Schütze, 2021), the quantity, relevance and even the order of the in-context examples (Brown et al., 2020a; Gao

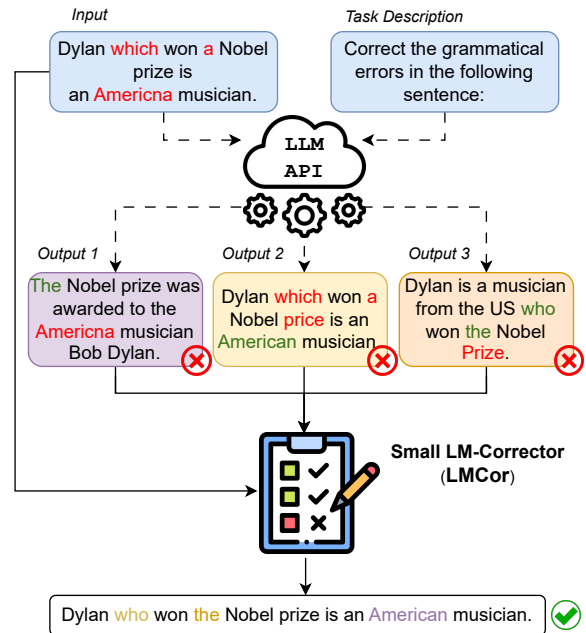


Figure 1: An illustration of our approach for grammatical error correction. We first prompt an LLM to generate multiple outputs via an API (dotted lines). Then we feed the generated candidates to the LM-corrector, a small model that is trained to rewrite them in order to generate the target sentence (solid lines).

et al., 2021; Liu et al., 2022; Zhang et al., 2023a; Lu et al., 2022). As a result, in-context learning often requires labour-intensive prompt engineering which does not always guarantee improved performance (Jiang et al., 2021).

Fine-tuning, on the other hand, has been proven highly effective when task-specific datasets are available, with smaller, fine-tuned models outperforming few-shot-prompted LLMs on various tasks (Lester et al., 2021; Chowdhery et al., 2022; Xu et al., 2023). While LLMs can also be fine-tuned to enhance their performance in specific tasks, there are several limitations. Firstly, the fine-tuning process can negatively impact the few-shot performance of LLMs on other tasks, leading to a trade-off between versatility and performance (Fu

*Research conducted during an internship at Google.

et al., 2023). Secondly, the increasing scale of LLMs makes fine-tuning on standard hardware computationally infeasible. To address these issues, parameter-efficient fine-tuning methods have been proposed (Houlsby et al., 2019; Lester et al., 2021; Li and Liang, 2021; Hu et al., 2022). Although these methods are more computationally efficient, they still require access to the model weights and substantial computational resources for loading and updating the model. Furthermore, due to the commercialization of LLMs, they are often available only through restricted inference APIs.

In light of these challenges, we propose a method that leverages only the outputs of LLMs to enhance their performance. Our work targets scenarios where training data is available, but extreme computational resources are not. To this end, we introduce LM-Corrector (LMCOR), a compact model that corrects the predictions produced by the LLM. Unlike fine-tuning methods, our approach operates directly on the LLM outputs, bypassing the need for access to their weights.

LMCOR capitalizes on the observation that LLMs can generate a diverse array of candidates for a single input which are often complimentary. Thus, it is possible to produce a superior output by optimally combining spans from different candidates (see Figure 2). LMCOR receives multiple candidates for a single input and learns to optimally rank, combine, and edit them, ultimately yielding more precise and higher-quality outputs. Figure 1 illustrates our approach, where LMCOR rewrites the first output of the LLM while incorporating correct spans from the second (*American*) and the third outputs (*the Nobel*) to produce the final, corrected output.

Our contributions can be summarized as follows. (1) We introduce LMCOR, a method to improve the performance of LLMs in the presence of training data without access to the model weights. (2) We conduct experiments on four natural language generation tasks where LLMs underperform specialized models. We demonstrate that a small LMCOR model with only 250 million parameters improves the performance of an LLM with 62 billion parameters, matching or even outperforming task-specific models. (3) We showcase that the corrector is robust to different prompts, alleviating the need for extensive prompt engineering. (4) We demonstrate the versatility of our approach showing that a single corrector can be effortlessly applied to different

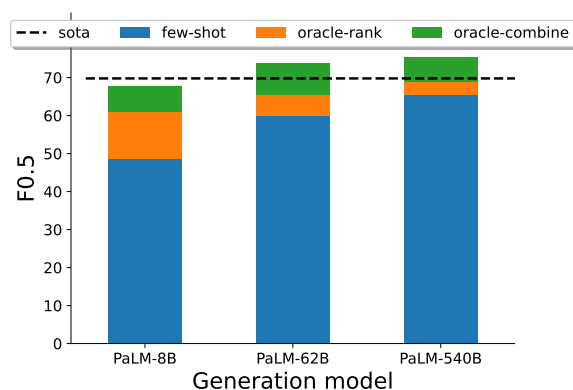


Figure 2: Potential of ranking (oracle-rank) and combining (oracle-combine) sampled candidates ($k=10$) from PaLM models of different scales for GEC.

LLMs as a plug-and-play module during inference. We make our code publicly available¹.

2 Correcting the Outputs of LLMs

In this section we present our computationally efficient approach that utilizes a small model, LMCOR, to correct the predictions of an LLM for a specific task. Unlike traditional fine-tuning methods, our approach does not require access to the weights of the LLM. Instead, as seen in Figure 1, we interact with the LLM only through an API, as is the case for some state-of-the-art commercial LLMs.

Headroom analysis Our approach is based on the insight that LLMs can generate a diverse pool of candidates for each input, with complementary strengths and weaknesses. Thus, an improved output can be produced by combining the correct parts of the corresponding candidates. To illustrate this, we experiment on the task of grammatical error correction (GEC) (Ng et al., 2014) using PaLM models (Chowdhery et al., 2022) of varying size, depicted in Figure 2. First, we observe that the few-shot PaLM models underperform fine-tuned 11B-parameter state-of-the-art (sota) GEC model (Rothe et al., 2021). However, by sampling 10 times from the LLM and employing an oracle to rank the samples (oracle-rank) or to combine correct spans (oracle-combine²), we obtain significant improvements, surpassing state-of-the-art.

This finding highlights the potential of leveraging multiple generations through ranking or com-

¹<https://github.com/GeorgeVern/lmcor>

²For the oracle-combine we compute the differing spans between the candidates and for each span we choose the one that has the smallest edit distance with the target.

binations to enhance the performance of the LLM via task-specific training. Motivated by this, we employ a smaller model, LMCOR, to predict the target given the original input and multiple candidates provided by the LLM.

Generating the candidates To start our pipeline, we first generate predictions from the LLM via in-context learning. Specifically, we prompt the model with a source sequence x , a verbal description of the task d , and a handful of demonstrations e , depicted as dashed lines in Figure 1. By sampling from the LLM with a temperature we obtain a diverse set of k candidates, $C = \{c_1, \dots, c_k\}$:

$$c_i \sim p_{LLM}(c|x, d, e) \quad \forall i = 1, 2, \dots, k. \quad (1)$$

Correcting the candidates Next, we feed the set of candidates along with the input sequence x to the corrector³ to generate the final, refined output.

$$\hat{y} = \arg \max_y p_{LMCOR}(y|x, C) \quad (2)$$

In order to train the corrector we fine-tune a small LM on the task-specific dataset augmented with candidates sampled from the LLM. Through this process, LMCOR learns to select the most promising among the generated outputs, combine different candidates and even make necessary edits to compose the desired target sentence. As we show in the following sections, even a small corrector can substantially improve the quality of LLM outputs, outperform standard fine-tuning, and reduce LLM sensitivity to different prompts.

3 Experiments & Results

3.1 Datasets and Models

We evaluate LMCOR on four natural language generation tasks: grammatical error correction on CoNLL-14 (Ng et al., 2014), data-to-text generation on E2E NLG (Novikova et al., 2017), summarization on XSum (Narayan et al., 2018) and machine translation on the English to German translation task from WMT22 (Kocmi et al., 2022).

In most of our experiments, we use the 62B version of PaLM (Chowdhery et al., 2022) as our large LM except for Section 4.2 where we vary the size of the LLM up to 540B parameters. For the machine translation task we use the 2.9B version of XGLM (Lin et al., 2022) as our LLM, since at the

³To indicate end-of-sequence boundaries for the input and the candidates, we use a sentinel token: $x[s]c_1[s]c_2[s] \dots [s]c_k$.

time of running this experiment, it was more easily accessible to the first author. We prompt the LLM with a task description and a number of demonstrations randomly selected from the respective validation set. We sample $k = 4$ times from the LLM with a temperature of 0.7 for PaLM and employ nucleus sampling (Holtzman et al., 2020) with $p = 0.6$ and a temperature of 0.6 for XGLM. Additionally, we include the greedy-decoded output as a candidate since initial results showed that it improves performance. We use T5.1.1⁴ base (Raffel et al., 2020) (250M parameters) as our model both for the LMCOR and the standard fine-tuning baseline. We choose the model based on the performance on the validation set. The outputs of the corrector and the T5 baseline are generated via beam search with a beam of size 5.

We compare our approach, LMCOR, with the following baselines: 1) in-context learning using the LLM (ICL), prompted with the same number of demonstrations, and 2) standard fine-tuning with a T5-base and PaLM. We also provide scores for 3) the reranking approach of Suzgun et al. (2022a) where they use Minimum Bayes Risk Decoding (MBRD) combined with an alignment function combined to select one among the candidates produced from the LLM. We use the same pool of candidates that are used as input to the corrector and employ Sim-LCS, a lexical similarity function based on longest common subsequence which achieved the best results across tasks among the alignment functions. We additionally provide the scores of an oracle reranker that selects the candidate with the smallest edit distance compared to the target as an upper-bound of reranking methods. Finally, we provide the results for a version of our approach that feeds the corrector with only the greedy-decoded candidate (single).

3.2 Grammatical Error Correction

Grammatical Error Correction (GEC) is a text-to-text task that requires correcting the grammatical errors while applying minimal changes to the original input sentence. Despite being trained on vast amounts of text, LLMs have been demonstrated to underperform task-specific models in this task (Yasunaga et al., 2021; Suzgun et al., 2022b).

We use the CoNLL-14 (Ng et al., 2014) dataset as our testset. Following previous

⁴https://github.com/google-research/text-to-text-transfer-transformer/blob/main/released_checkpoints.md

Model	F0.5
T5-base	59.38
PaLM-62B (ICL)	59.92
+ MBRD-Sim-LCS	58.87
+ Oracle Reranker	63.88
+ LMCOR (single)	<u>62.47</u>
+ LMCOR (mult.)	62.48

Table 1: Results of our approach in GEC (CoNLL-14). The first group indicates fine-tuned models, the second group in-context learning with reranking and the final group provides the scores for LMCOR. The best scores are in **bold** and the second best ones are underlined.




Figure 3: The effect of dataset size for standard fine-tuning and LMCOR. Results are reported on GEC.

work (Rothe et al., 2021), we use the combination of the FCE (Yannakoudakis et al., 2011) and W&I (Bryant et al., 2019) datasets (60k examples) for training and validation. We report $F_{0.5}$ scores obtained with the MaxMatch scorer (Dahlmeier and Ng, 2012)⁵. We use 5 demonstrations in the LLM prompt in order to generate the candidates during training and inference.

The results presented in Table 1 show that standard fine-tuning and in-context learning exhibit comparable performance in GEC while our approach significantly outperforms both, by 3 and 2.5 $F_{0.5}$ points respectively. It is worth noting that MBRD does not yield any improvement over standard few-shot prompting. However, as expected the use of an oracle to rank the produced hypotheses results in a considerable performance boost. Although LMCOR does not surpass the performance of the oracle it manages to significantly close the gap, demonstrating the ability of the corrector to identify high-quality candidates from the model.

Additionally, we compare the performance of standard fine-tuning and LMCOR across varying numbers of training instances. As illustrated in

⁵<https://www.comp.nus.edu.sg/nlp/conll14st.html>

Model	R-2	R-L
T5-base	45.3	52.8
PaLM-62B* (FT)	45.2	–
PaLM-540B* (FT)	<u>45.3</u>	52.3
PaLM-62B (ICL)	35.1	45.6
+ MBRD-Sim-LCS	35.7	46.2
+ Oracle Reranker	37.1	50.4
+ LMCOR (single)	44.8	<u>52.8</u>
+ LMCOR (mult.)	45.6	53.4

Table 2: Results of our approach in E2E NLG (cleaned). Results with * are reported from the original paper (Chowdhery et al., 2022). The first group indicates fine-tuned models, the second group in-context learning with reranking and the final group provides the scores for LMCOR. The best scores are in **bold** and the second best ones are underlined.

Figure 3 our approach consistently outperforms the baseline for all dataset sizes. The gap is particularly pronounced when the training dataset is limited, consisting of only 1k examples, resulting in a substantial difference of 15 points in $F_{0.5}$. The sample efficiency of LMCOR can be attributed to its ability to leverage the candidates generated by the LLM to produce more accurate outputs. We note that in this low-resource scenario, both trained models perform worse than few-shot prompting. This outcome is expected as the extensive pretraining of LLMs on language generation enables them to perform grammatical error correction out of the box. As dataset increases to 10k examples we observe that LMCOR performs on par with the LLM while the baseline continues to underperform. Beyond this threshold, LMCOR surpasses both in-context learning and fine-tuning by utilizing both the training data and the candidates.

3.3 Data-to-text

The next task we evaluate on is E2E NLG (Novikova et al., 2017), a data-to-text task where the input is a number of key-value pairs about a restaurant and the output is a short description of the restaurant in natural language. We use the cleaned version of the dataset, E2E NLG (cleaned) (Dušek et al., 2019) and the default splits for training (35k examples), validation and testing. We use 5 demonstrations to produce the candidates for the corrector both during training and inference. We report ROUGE-2 and ROUGE-L (Lin, 2004) scores.

Table 2 presents a comparison between standard fine-tuning with a T5-base model, in-context learn-

ing, and fine-tuning using PaLM models. Notably, standard fine-tuning with a T5-base significantly outperforms in-context learning and achieves results comparable to fine-tuning with the much larger PaLM models⁶. While reranking techniques improve the performance of few-shot prompting, even the oracle approach falls short of the performance achieved by the fine-tuned T5-base model. This highlights the primary limitation of reranking approaches, particularly for challenging tasks for LLMs, where their upper bound solely depends on the quality of candidates.

In contrast, the performance of LMCOR does not exclusively rely on the quality of the candidates. The corrector module has the ability to edit the LLM-generated candidates, leading to more accurate outputs. As a result, LMCOR demonstrates the best overall performance for the E2E task, surpassing even the fine-tuned PaLM-540B model by 1 point in ROUGE-L. An important characteristic of our approach is the ability of the corrector to observe multiple candidates for a single input. This enables LMCOR to combine candidates in order to compose a more refined answer. This is supported by the performance discrepancy between LMCOR (single) and LMCOR (mult.) highlighting the effectiveness of leveraging multiple candidates.

3.4 Summarization

The third task that we consider is abstractive summarization. Specifically, we use XSum (Narayan et al., 2018) with the default train (204k examples), validation and test splits. Due to the length of the articles we truncate the inputs and only use 1 demonstration to prompt the LLM. To handle the increased sequence length in the input of the corrector we again truncate the articles and use a maximum sequence length of 2048 tokens. We report ROUGE-1, ROUGE-2 and ROUGE-L scores.

The results of Table 3 reveal that standard fine-tuning outperforms in-context learning for the XSum dataset. Specifically, the fine-tuned T5 and PaLM-62B models outperform in-context learning by 15 and 18 points in ROUGE-2 respectively. The difficulty of the task, which involves summarizing an article into a single sentence, poses challenges for in-context learning, while the substantial dataset size of 204k examples favors fine-tuning. However, the use of a corrector module leads to significant

⁶Chowdhery et al. (2022) attribute the mediocre performance of the fine-tuned PaLM models to the small dataset size and the ‘significant mismatch with the pre-training corpus’.

Model	R-1	R-2	R-L
T5-base	38.64	16.98	31.41
PaLM-62B* (FT)	–	18.5	–
PaLM-540B* (FT)	–	21.2	36.5
PaLM-62B (ICL)	28.18	10.50	22.38
PaLM-540B (ICL)	29.88	11.75	23.83
+ LMCOR (single)	36.98	16.41	30.20
+ LMCOR (mult.)	<u>37.62</u>	<u>16.50</u>	<u>30.67</u>

Table 3: Results of our approach on XSum. Results with * are reported from the original paper (Chowdhery et al., 2022). The first group indicates fine-tuned models, the second group in-context learning and the final group provides the scores for LMCOR. The best scores are in **bold** and the second best ones are underlined.

improvements over in-context learning with the 62B PaLM model, resulting in performance gains of 6 points in ROUGE-2 and 8 points in ROUGE-L. Notably, LMCOR outperforms few-shot learning with the largest 540B PaLM model. Again, the use of multiple candidates by the corrector further enhances the performance of LMCOR confirming our hypothesis regarding the complementarity of the outputs generated by the LLM.

On the other hand, the LMCOR performs slightly worse than standard fine-tuning. This result is somewhat unexpected since the input for the corrector is a strict superset of the model input in the fine-tuning setting. We attribute the slight drop in performance to the poor quality of the provided candidates, which introduces undesirable noise to the corrector’s input, a hypothesis that we test in Section 4.3. Additionally, the long-range dependencies that are introduced by simultaneously processing the article and the generated summaries might also contribute to the performance gap between LMCOR and standard fine-tuning.

3.5 Machine Translation

The final task in our evaluation is machine translation (MT). For this task we use the English to German language pair from WMT22 (Kocmi et al., 2022) as our test set and the corresponding pair from WMT21 (Akhbardeh et al., 2021) as our validation set. Our training data consists of 200k examples sampled from the News Commentary v16 corpus⁷. During both training and inference, we prompt the LLM with 5 demonstrations to generate the candidates. We report scores using traditional surface-based MT evaluation metrics like

⁷<https://www.statmt.org/wmt22/translation-task.html>

Model	BLEU	COMET	BLEURT
T5-base	23.32	75.22	64.57
XGLM-2.9B (ICL)	17.32	74.54	66.47
+ MBRD-Sim-CLS	18.01	74.82	66.73
+ Oracle Reranker	21.21	75.55	66.90
+ LMCOR (single)	<u>24.51</u>	<u>76.81</u>	<u>67.23</u>
+ LMCOR (mult.)	25.15	77.45	68.41

Table 4: Results of our approach on WMT22 En->De. The first group indicates fine-tuned models, the second group in-context learning and the final group provides the scores for LMCOR. The best scores are in **bold** and the second best ones are underlined.

BLEU (Papineni et al., 2002), as well as more recent neural-based metrics such as COMET-22 (Rei et al., 2022) and BLEURT (Sellam et al., 2020).

The findings presented in Table 4 indicate that, similar to previous tasks, standard fine-tuning outperforms in-context learning for MT across two of the three considered metrics. While the minimum Bayes risk reranking approach shows some improvements, it fails to significantly narrow the gap with the fine-tuned baseline. Notably, the oracle reranker manages to surpass T5 in terms of COMET scores, although T5 still outperforms in terms of BLEU.

Our proposed approach achieves substantial gains, surpassing both fine-tuning and in-context learning, as well as the reranking approaches, including the oracle. This observation once again highlights the limitations of reranking approaches, especially when dealing with low-quality candidates. Specifically, the version of LMCOR using a single candidate demonstrates improvements of 4.5 BLEU points, 2 points in terms of COMET, and 0.5 points in terms of BLEURT compared to the best scores achieved either by fine-tuning or reranking. The inclusion of multiple candidates yields further gains ranging from 0.6 to 1.2 points, depending on the metric, underscoring the benefits of a diverse candidate pool.

4 Robustness Analysis

4.1 Different prompts

We have demonstrated that LMCOR enhances the performance of LLMs by refining their generated predictions. However, the few-shot paradigm remains appealing since it does not require training or access to a dataset, except for a handful of examples used in the prompt. While prompt construction is not computationally intensive, there is no consensus on the optimal selection, number and even

Model	set 1	set 2	set 3	mean	std
PaLM-62B	59.9	58.9	56.2	58.6	1.9
+ LMCOR	62.5	62.3	62.9	62.6	0.3

Table 5: Mean $F_{0.5}$ score and standard deviation (std) using different sets of demonstrations for in-context learning vs. our approach for GEC.

order of examples in the prompt which can lead to variations in predictions and significantly impact LLM performance (Lu et al., 2022). To investigate whether LMCOR is similarly affected by this variability we use three different sets of 5 demonstrations to prompt the LLM for the task of GEC. We then feed the generated predictions as input to the corrector. It is important to note that we trained LMCOR only once using the candidates generated with the original set of demonstrations (set 1) and simply swapped candidates during inference.

Table 5 highlights the significant variance in LLM performance depending on the selection of demonstrations with a difference of 3.7 between the highest and the lowest $F_{0.5}$ score. In contrast, LMCOR remains unaffected and achieves competitive performance even when the quality of candidates significantly deteriorates (set 3). This demonstrates the capability of LMCOR to compensate for candidates of poor quality by performing edits on them. The robustness of the corrector against prompts of varying quality (with a variance of 0.3 compared to 1.9 for the LLM) suggests that it can mitigate the need for extensive prompt engineering.

4.2 Different LLMs

In the previous experiment we demonstrated the robustness of the corrector against candidates of varying quality. In this set of experiments we further examine the robustness of our approach by testing whether a corrector can be used interchangeably with different LMs without retraining. It is important to note that we only trained the corrector once and performed inference by swapping the LLM responsible for generating the candidates.

Initially, we focus on LMs from the same family of models, namely PaLM (Chowdhery et al., 2022), which share similar architectures and training data but differ in the number of parameters. Table 6 presents the results of applying the corrector to different PaLM models than the one it was originally trained on (62B). Across all scales, the LMCOR consistently outperforms standard fine-tuning and in-context learning, with the exception of the 540B

T5-base	59.38		
PaLM (ICL)	8B	62B	540B
	48.62	59.92	65.37
+ LMCOR (single)	61.40	62.48	63.55
+ LMCOR (mult.)	61.89	62.47	65.16

Table 6: Results ($F_{0.5}$) of applying the corrector to LMs of different scale during inference for GEC.

Model	R-2	R-L
GPT3-Codex (ICL)*	34.2	44.4
+ MBRD-BLEURT*	36.4	46.5
+ LMCOR (mult.)	44.8	53.0

Table 7: Applying the corrector to different family of LMs during inference for E2E NLG. *: the results as reported in the original paper.

model where the performance is comparable. The corrector achieves significant gains when applied to the 8B PaLM model, with an improvement of +13 points in $F_{0.5}$. This further highlights the ability of LMCOR to compensate for low-quality candidates by merging and correcting them in order to achieve competitive performance. Furthermore, we observe that using a single candidate for LMCOR leads to inferior performance in all cases, except for the 62B PaLM model. This finding suggests that the existence of diverse candidates prevents the model from overfitting to the outputs of a specific LLM, thereby enhancing its generalization capabilities.

As a next step, we explore the application of LMCOR to an LLM from a distinct family of models, specifically Codex⁸ which is a GPT3-like model trained on code (Chen et al., 2021). To compare the effectiveness of LMCOR with MRBD reranking we utilize BLEURT (Sellam et al., 2020) as the alignment function, as it has been reported to achieve the highest scores for the E2E NLG task (Suzgun et al., 2022b). It is important to note that while the reranking approach samples 16 outputs from the LLM, we only use 5 for the corrector. The results in Table 7 demonstrate the superior performance of LMCOR over reranking. In particular we observe a performance boost of 10 points in ROUGE-2 when compared to in-context learning, whereas MBRD achieves a mere 2-point improvement.

The previous findings highlight the remarkable out-of-domain robustness of LMCOR and its ability to seamlessly integrate with various LLMs, as a versatile solution for enhancing their performance.

⁸We use the outputs of code-davinci-002 provided by Suzgun et al. (2022a).

Model	R-1	R-2	R-L	BLEU
Pegasus (FT)	45.48	23.88	38.18	16.72
+ LMCOR	45.76	23.78	38.28	17.00

Table 8: Applying the corrector to state-of-the-art summarization model. Results are reported on XSum.

This versatility not only holds promise for applying a single corrector to multiple LLMs but also for training correctors with future, more capable LLMs

4.3 Task-specific models

To further assess the versatility of LMCOR we extend our investigation to specialized models. Specifically, we train a corrector using candidates produced by Pegasus (Zhang et al., 2020), a state-of-the-art summarization model, via beam search. The results of Table 8 reveal that LMCOR provides gains even for models that have undergone pre-training and fine-tuning tailored to the task. Although the gains are relatively modest compared to PaLM, this discrepancy can be attributed to the already high performance of Pegasus and the lack of diversity of the beam-generated candidates, which is essential for the corrector. The increased performance of the corrector when applied to Pegasus compared to PaLM supports our intuition regarding the noise introduced by low-quality candidates.

5 Analysis

5.1 Importance of the source

The input of the corrector consists of the source sentence and a number of candidates generated by the LLM (Equation 2). In previous sections we demonstrated that the use of multiple candidates improves in-domain performance and out-of-domain robustness. In this experiment, we focus on the importance of the source sentence to LMCOR. To examine this, we train a corrector that receives only the candidates as input, without access to the source. The results for E2E NLG, presented in Table 9 reveal a noticeable decline in performance when the source sentence is removed. This decrease can be attributed to the inability of the corrector to produce outputs that are faithful to the input in the absence of the source sentence.

5.2 Scaling the Corrector

We showed that a corrector with 250M parameters, can effectively refine the predictions of LLMs for specific tasks. This raises the question: *is training*

Model	R-2	R-L
PaLM-62B (ICL)	35.1	45.6
+ LMCOR	45.6	53.4
- source sentence	44.5	53.1

Table 9: The importance of the source sentence for the corrector. Results are reported on E2E NLG (clean).




Figure 4: The effect of scaling for LMCOR and fine-tuning. Results are reported on GEC.

a corrector still valuable if we have the computational resources to train a very large model? To investigate this we train the largest version of T5, T5-xxl with 11B parameters, both through standard fine-tuning and as a corrector for GEC. We note that, in this scenario, the sizes of the fine-tuned model and the LLM are comparable (11B vs. 62B).

As shown in Figure 4 both LMCOR and the fine-tuned T5 benefit from scaling, exhibiting higher $F0.5$ scores as their parameter count increases from 250 million to 11 billion. At the 11 billion scale, both models significantly outperform in-context learning with the 62B PaLM model. The corrector continues to outperform the baseline as model scale increases, although the gap in performance narrows from 3 to 0.5 points in $F0.5$. We attribute this reduction to the enhanced competence of the larger T5 model. At the scale of 11 billion parameters, the scores obtained by the T5 model alone surpass those obtained by the LLM, indicating that the LLM-generated outputs are of lower quality compared to the ones that the model would produce independently. Therefore, while the additional input to the model remains beneficial, its impact diminishes to some extent, considering the performance disparity between T5 and PaLM.

6 Related Work

Since the introduction of in-context learning, prior research has primarily focused on improving the

few-shot performance of LLMs. One approach suggests prompting the model to generate rationales or chain-of-thoughts (Nye et al., 2021; Wei et al., 2022; Kojima et al., 2022) or to decompose the problem into simpler ones (Press et al., 2022; Zhou et al., 2023; Pilault et al., 2023), simulating a reasoning process prior to generating the answer. These prompting techniques are complementary to our approach and can provide improved candidates for the corrector.

Another strategy to improve the performance of LLMs without training is reranking, i.e. selecting the most promising from a pool of candidates obtained by sampling from the model. Reranking approaches include training different models as the ranker (Cobbe et al., 2021), using task-specific ranking functions (Suzgun et al., 2022b; Fernandes et al., 2022), majority voting (Wang et al., 2023) or minimum Bayes risk decoding (MBRD) (Suzgun et al., 2022a; Freitag et al., 2022). Although these approaches can improve the few-shot performance of an LLM, they are upper-bounded by the quality of the generated candidates.

While fine-tuning large LLMs can enhance performance, the substantial computational requirements have prompted the development of parameter-efficient fine-tuning methods (PEFT) (He et al., 2022). These approaches introduce a small number of additional parameters (relative to the full model) to be trained while the rest of the model is frozen. The newly-added parameters can come in the form of embeddings that are appended to the encoded sequence (Li and Liang, 2021; Lester et al., 2021), MLPs that are added in-between layers, namely adapters (Houlsby et al., 2019; Karimi Mahabadi et al., 2021) or rank-decomposition matrices that are added in parallel to the existing layers (Hu et al., 2022; Zhang et al., 2023b). Although these works decrease the computational load of fine-tuning they still require loading and backpropagating through the model, which can be prohibitive for LLMs. Our work shares the same motivation with PEFT methods, with the introduced parameters being essentially another, smaller model. However, our method does not have the memory requirements of PEFT since the corrector operates directly on the model’s outputs and does not require access to the model’s weights.

An alternative line of work proposes providing feedback to the LLM in order to revise and enhance its predictions. The feedback can be obtained

from external models such as Google Search, document retrievers, compilers (Gao et al., 2021; Yao et al., 2023; Peng et al., 2023; Gou et al., 2023), or from a separate model trained to provide feedback on LLM outputs with additional supervision (Paul et al., 2023; Peng et al., 2023; Akyürek et al., 2023). While leveraging the LLM itself to generate feedback has been explored (Madaan et al., 2023; Shinn et al., 2023), it tends to yield lower-quality feedback (Akyürek et al., 2023; Gou et al., 2023; Huang et al., 2023) and involves multiple passes and extensive prompt engineering for each LLM operation. In contrast, our approach is task-agnostic and requires a single pass from the LLM, with little to no prompt engineering, offering an efficient solution for enhancing LLM outputs.

Recently, studies have highlighted the potential of smaller, task-specific models to complement the predictions of an LLM. Xu et al. (2023) explore a framework where candidates produced by task-specific models are fed to an LLM, primarily targeting classification task while LMCOR is better-suited for open-ended generation tasks. Welleck et al. (2023) train a smaller model to iteratively improve sequences generated by LLMs. In contrast to our method, they rely on unlabeled data and sample extensively from the LLM to obtain a large pool of candidates. They assume the availability of a value function that assigns scores to each candidate and create input-output pairs by sorting candidates based on their scores. Unlike their approach, we demonstrate that a compact corrector can perform effectively across various tasks. Additionally, our approach is more efficient during inference, since the ability of LMCOR to process multiple candidates simultaneously eliminates the need for multiple passes.

Concurrently, researchers have begun to leverage the complementary nature of LLM-generated outputs during inference. Farinhas et al. (2023) use an LLM to combine its generated outputs for machine translation, although they find that reranking methods incorporating external modules, such as quality estimation metrics (Zerva et al., 2022), prove to be more efficient. Meanwhile, Vernikos and Popescu-Belis (2024) propose an approach that uses a quality estimation metric to combine the outputs of LLMs or MT models. Similar to our method, they exploit the diversity of LLM outputs by identifying divergent spans among candidates and merging them based on the metric.

Most relevant to our approach is the work by Jiang et al. (2023) where they propose a method to ensemble LLMs. Their pipeline consists of i) sampling a large pool of candidates, ii) selecting top candidates via multiple pairwise comparisons through a trained reranker and iii) fusing them using a similar technique as LMCOR. While our approach could be extended to multiple LLMs we demonstrate improvements with a single LLM, leveraging the complementarity of the generations. In addition, our approach is more efficient since we use a much smaller model as the corrector (3B vs 250M) and do not introduce additional training and inference steps for ranking the outputs.

7 Conclusion

In this work, we introduce LMCOR, a novel approach that leverages a small corrector module to enhance the performance of LLMs in the presence of training data. LMCOR leverages the diversity of the LLM generations to rank, edit and combine the candidates. Unlike parameter-efficient fine-tuning methods, our approach does not require access to the model or substantial computational resources. Our experiments demonstrate that even a relatively small corrector (250M) can improve the performance of a much larger LM (62B), while exhibiting robustness against different prompts. Furthermore, we showcase that the corrector can be successfully applied to models of different scale or architecture without any retraining. These findings offer a promising solution for improving LLM performance in a practical and resource-efficient manner and open up new possibilities for the utilization and deployment of LLMs in real-world applications alongside smaller task-specific models.

Acknowledgments

We are grateful for their support to the Swiss National Science Foundation (DOMAT grant n. 175693, On-demand Knowledge for Document-level Machine Translation), and to the Institute for ICT at HEIG-VD. We thank Andrei Popescu-Belis and Katerina Margatina for their valuable comments and fruitful discussions.

Limitations

Additional latency While our approach enhances the performance of LLMs on the considered tasks, it also introduces additional latency. Instead of a single pass from the LLM our pipeline involves

sampling multiple candidates in parallel and performing an additional inference step using a much smaller model. Although the additional latency is small, it could be critical for low-latency applications that require real-time responses.

Strong fine-tuned models While our approach demonstrates gains over in-context learning and reranking it may not always achieve the same level of performance as fine-tuning approaches. Our results in XSum and Figure 3 indicate that fine-tuning remains a powerful method for smaller models when ample data is available. Additionally, scaling fine-tuned models instead of using an off-the-shelf LLM might be a better alternative in certain cases, as discussed in Section 5.2.

Additional tasks and models Due to time and budget limit our experiments cover 4 natural language generation tasks and they could be extended to other kinds of tasks such as reasoning. Additionally while LMCOR shows promising results when combined with different LLMs, even during inference, it would be interesting to apply our approach to a broader selection of LLMs that are currently available.

Human evaluation We agree that automatic metrics have limitations. While the selection of metrics aligns with prior work, human evaluation could provide us with more reliable and comprehensive evaluation results. However, due to the number of models and the amount of generation candidates, we could not afford large-scale human evaluation.

Broader Impact

Given that the proposed approach combines LLMs and small corrector models for improved performance, it is important to acknowledge that it shares the potential social biases associated with LLMs. While our work focuses on improving the predictions of LLMs on specific datasets rather than open-ended generation, it is improbable that our approach amplifies these biases to a greater extent than other methods. Nonetheless, it is important to investigate whether LMCOR has any impact on either amplifying or mitigating these biases.

References

Farhad Akhbardeh, Arkady Arkhangorodsky, Magdalena Biesialska, Ondřej Bojar, Rajen Chatterjee, Vishrav Chaudhary, Marta R. Costa-jussa,

Cristina España-Bonet, Angela Fan, Christian Federmann, Markus Freitag, Yvette Graham, Roman Grundkiewicz, Barry Haddow, Leonie Harter, Kenneth Heafield, Christopher Homan, Matthias Huck, Kwabena Amponsah-Kaakyire, Jungo Kasai, Daniel Khashabi, Kevin Knight, Tom Kocmi, Philipp Koehn, Nicholas Lourie, Christof Monz, Makoto Morishita, Masaaki Nagata, Ajay Nagesh, Toshiaki Nakazawa, Matteo Negri, Santanu Pal, Allahsera Auguste Tapo, Marco Turchi, Valentin Vydrin, and Marcos Zampieri. 2021. [Findings of the 2021 conference on machine translation \(WMT21\)](#). In *Proceedings of the Sixth Conference on Machine Translation*, pages 1–88.

Afra Feyza Akyürek, Ekin Akyürek, Aman Madaan, Ashwin Kalyan, Peter Clark, Derry Wijaya, and Niket Tandon. 2023. [R14f: Generating natural language feedback with reinforcement learning for repairing model outputs](#).

Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020a. [Language models are few-shot learners](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901.

Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020b. [Language models are few-shot learners](#). *Advances in neural information processing systems*, 33:1877–1901.

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.

Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde, Jared Kaplan, Harrison Edwards, Yura Burda, Nicholas Joseph, Greg Brockman, Alex Ray, Raul Puri, Gretchen Krueger, Michael Petrov, Heidy Khlaaf, Girish Sastry, Pamela Mishkin, Brooke Chan, Scott Gray, Nick Ryder, Mikhail Pavlov, Alethea Power, Lukasz Kaiser, Mohammad Bavarian, Clemens Winter, Philippe Tillet, Felipe Petroski Such, David W. Cummings, Matthias Plappert, Fotios Chantzis, Elizabeth Barnes, Ariel Herbert-Voss, William H. Guss, Alex Nichol, Igor Babuschkin, S. Arun Balaji, Shantanu Jain, Andrew Carr, Jan Leike, Joshua Achiam, Vedant Misra, Evan Morikawa, Alec Radford, Matthew M. Knight, Miles Brundage, Mira Murati, Katie Mayer, Peter Welinder, Bob McGrew, Dario Amodei, Sam McCandlish, Ilya

- Sutskever, and Wojciech Zaremba. 2021. Evaluating large language models trained on code. *ArXiv*, abs/2107.03374.
- Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, Parker Schuh, Kensen Shi, Sasha Tsvyashchenko, Joshua Maynez, Abhishek Rao, Parker Barnes, Yi Tay, Noam Shazeer, Vinodkumar Prabhakaran, Emily Reif, Nan Du, Ben Hutchinson, Reiner Pope, James Bradbury, Jacob Austin, Michael Isard, Guy Gur-Ari, Pengcheng Yin, Toju Duke, Anselm Levskaya, Sanjay Ghemawat, Sunipa Dev, Henryk Michalewski, Xavier Garcia, Vedant Misra, Kevin Robinson, Liam Fedus, Denny Zhou, Daphne Ippolito, David Luan, Hyeontaek Lim, Barret Zoph, Alexander Spiridonov, Ryan Sepassi, David Dohan, Shivani Agrawal, Mark Omernick, Andrew M. Dai, Thanumalayan Sankaranarayanan Pillai, Marie Pellat, Aitor Lewkowycz, Erica Moreira, Rewon Child, Oleksandr Polozov, Katherine Lee, Zongwei Zhou, Xuezhi Wang, Brennan Saeta, Mark Diaz, Orhan Firat, Michele Catasta, Jason Wei, Kathy Meier-Hellstern, Douglas Eck, Jeff Dean, Slav Petrov, and Noah Fiedel. 2022. [Palm: Scaling language modeling with pathways](#).
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. 2021. [Training verifiers to solve math word problems](#). *CoRR*, abs/2110.14168.
- 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.
- Ondřej Dušek, David M. Howcroft, and Verena Rieser. 2019. [Semantic noise matters for neural natural language generation](#). In *Proceedings of the 12th International Conference on Natural Language Generation*, pages 421–426.
- António Farinhas, José de Souza, and Andre Martins. 2023. [An empirical study of translation hypothesis ensembling with large language models](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 11956–11970.
- Patrick Fernandes, António Farinhas, Ricardo Rei, José G. C. de Souza, Perez Ogayo, Graham Neubig, and Andre Martins. 2022. [Quality-aware decoding for neural machine translation](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1396–1412.
- Markus Freitag, David Grangier, Qijun Tan, and Bowen Liang. 2022. [High quality rather than high model probability: Minimum Bayes risk decoding with neural metrics](#). *Transactions of the Association for Computational Linguistics*, 10:811–825.
- Yao Fu, Hao Peng, Litu Ou, Ashish Sabharwal, and Tushar Khot. 2023. [Specializing smaller language models towards multi-step reasoning](#). In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 10421–10430.
- Tianyu Gao, Adam Fisch, and Danqi Chen. 2021. [Making pre-trained language models better few-shot learners](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 3816–3830.
- Zhibin Gou, Zhihong Shao, Yeyun Gong, Yelong Shen, Yujiu Yang, Nan Duan, and Weizhu Chen. 2023. [Critic: Large language models can self-correct with tool-interactive critiquing](#).
- Junxian He, Chunting Zhou, Xuezhe Ma, Taylor Berg-Kirkpatrick, and Graham Neubig. 2022. [Towards a unified view of parameter-efficient transfer learning](#). In *International Conference on Learning Representations*.
- Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. 2020. [The curious case of neural text de-generation](#). In *International Conference on Learning Representations*.
- Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin De Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. 2019. [Parameter-efficient transfer learning for NLP](#). In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 2790–2799.
- Edward J Hu, yelong shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2022. [LoRA: Low-rank adaptation of large language models](#). In *International Conference on Learning Representations*.
- Jie Huang, Xinyun Chen, Swaroop Mishra, Huaixiu Steven Zheng, Adams Wei Yu, Xinying Song, and Denny Zhou. 2023. [Large language models cannot self-correct reasoning yet](#).
- Dongfu Jiang, Xiang Ren, and Bill Yuchen Lin. 2023. [LLM-blender: Ensembling large language models with pairwise ranking and generative fusion](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 14165–14178.
- Zhengbao Jiang, Jun Araki, Haibo Ding, and Graham Neubig. 2021. [How can we know when language models know? on the calibration of language models for question answering](#). *Transactions of the Association for Computational Linguistics*, 9:962–977.
- Rabeeh Karimi Mahabadi, James Henderson, and Sebastian Ruder. 2021. [Compacter: Efficient low-rank hypercomplex adapter layers](#). *Advances in Neural Information Processing Systems*, 34:1022–1035.

- Tom Kocmi, Rachel Bawden, Ondřej Bojar, Anton Dvorkovich, Christian Federmann, Mark Fishel, Thamme Gowda, Yvette Graham, Roman Grundkiewicz, Barry Haddow, Rebecca Knowles, Philipp Koehn, Christof Monz, Makoto Morishita, Masaaki Nagata, Toshiaki Nakazawa, Michal Novák, Martin Popel, and Maja Popović. 2022. [Findings of the 2022 conference on machine translation \(WMT22\)](#). In *Proceedings of the Seventh Conference on Machine Translation (WMT)*, pages 1–45.
- Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. 2022. [Large language models are zero-shot reasoners](#). In *ICML 2022 Workshop on Knowledge Retrieval and Language Models*.
- Brian Lester, Rami Al-Rfou, and Noah Constant. 2021. [The power of scale for parameter-efficient prompt tuning](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 3045–3059.
- Xiang Lisa Li and Percy Liang. 2021. [Prefix-tuning: Optimizing continuous prompts for generation](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 4582–4597.
- Chin-Yew Lin. 2004. [ROUGE: A package for automatic evaluation of summaries](#). In *Text Summarization Branches Out*, pages 74–81.
- Xi Victoria Lin, Todor Mihaylov, Mikel Artetxe, Tianlu Wang, Shuohui Chen, Daniel Simig, Myle Ott, Nanman Goyal, Shruti Bhosale, Jingfei Du, Ramakanth Pasunuru, Sam Shleifer, Punit Singh Koura, Vishrav Chaudhary, Brian O’Horo, Jeff Wang, Luke Zettlemoyer, Zornitsa Kozareva, Mona Diab, Veselin Stoyanov, and Xian Li. 2022. [Few-shot learning with multilingual generative language models](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 9019–9052.
- Jiachang Liu, Dinghan Shen, Yizhe Zhang, Bill Dolan, Lawrence Carin, and Weizhu Chen. 2022. [What makes good in-context examples for GPT-3?](#) In *Proceedings of Deep Learning Inside Out (DeeLIO 2022): The 3rd Workshop on Knowledge Extraction and Integration for Deep Learning Architectures*, pages 100–114.
- Yao Lu, Max Bartolo, Alastair Moore, Sebastian Riedel, and Pontus Stenetorp. 2022. [Fantastically ordered prompts and where to find them: Overcoming few-shot prompt order sensitivity](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8086–8098.
- Aman Madaan, Niket Tandon, Prakhar Gupta, Skyler Hallinan, Luyu Gao, Sarah Wiegrefe, Uri Alon, Nouha Dziri, Shrimai Prabhunoye, Yiming Yang, Sean Welleck, Bodhisattwa Prasad Majumder, Shashank Gupta, Amir Yazdanbakhsh, and Peter Clark. 2023. [Self-refine: Iterative refinement with self-feedback](#).
- Shashi Narayan, Shay B Cohen, and Mirella Lapata. 2018. [Don’t give me the details, just the summary! topic-aware convolutional neural networks for extreme summarization](#). *arXiv preprint arXiv:1808.08745*.
- 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.
- Jekaterina Novikova, Ondřej Dušek, and Verena Rieser. 2017. [The E2E dataset: New challenges for end-to-end generation](#). In *Proceedings of the 18th Annual SIGdial Meeting on Discourse and Dialogue*, pages 201–206.
- Maxwell I. Nye, Anders Johan Andreassen, Guy Gur-Ari, Henryk Michalewski, Jacob Austin, David Bieber, David Dohan, Aitor Lewkowycz, Maarten Bosma, David Luan, Charles Sutton, and Augustus Odena. 2021. [Show your work: Scratchpads for intermediate computation with language models](#). *CoRR*, abs/2112.00114.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. [Bleu: a method for automatic evaluation of machine translation](#). In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318.
- Debjit Paul, Mete Ismayilzada, Maxime Peyrard, Beatriz Borges, Antoine Bosselut, Robert West, and Boi Faltings. 2023. [Refiner: Reasoning feedback on intermediate representations](#).
- Baolin Peng, Michel Galley, Pengcheng He, Hao Cheng, Yujia Xie, Yu Hu, Qiuyuan Huang, Lars Liden, Zhou Yu, Weizhu Chen, and Jianfeng Gao. 2023. [Check your facts and try again: Improving large language models with external knowledge and automated feedback](#).
- Jonathan Pilault, Xavier Garcia, Arthur Bražinskas, and Orhan Firat. 2023. [Interactive-chain-prompting: Ambiguity resolution for crosslingual conditional generation with interaction](#). *arXiv preprint arXiv:2301.10309*.
- Ofir Press, Muru Zhang, Sewon Min, Ludwig Schmidt, Noah A. Smith, and Mike Lewis. 2022. [Measuring and narrowing the compositionality gap in language models](#).
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. [Exploring the limits of transfer learning with a unified text-to-text transformer](#). *J. Mach. Learn. Res.*, 21(1).

- Ricardo Rei, José G. C. de Souza, Duarte Alves, Chrysoula Zerva, Ana C Farinha, Taisiya Glushkova, Alon Lavie, Luisa Coheur, and André F. T. Martins. 2022. [COMET-22: Unbabel-IST 2022 submission for the metrics shared task](#). In *Proceedings of the Seventh Conference on Machine Translation (WMT)*, pages 578–585.
- 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.
- Timo Schick and Hinrich Schütze. 2021. [It’s not just size that matters: Small language models are also few-shot learners](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2339–2352.
- Thibault Sellam, Dipanjan Das, and Ankur Parikh. 2020. [BLEURT: Learning robust metrics for text generation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7881–7892.
- Taylor Shin, Yasaman Razeghi, Robert L. Logan IV, Eric Wallace, and Sameer Singh. 2020. [AutoPrompt: Eliciting Knowledge from Language Models with Automatically Generated Prompts](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4222–4235.
- Noah Shinn, Federico Cassano, Beck Labash, Ashwin Gopinath, Karthik Narasimhan, and Shunyu Yao. 2023. [Reflexion: Language agents with verbal reinforcement learning](#).
- Mirac Suzgun, Luke Melas-Kyriazi, and Dan Jurafsky. 2022a. [Follow the wisdom of the crowd: Effective text generation via minimum bayes risk decoding](#).
- Mirac Suzgun, Luke Melas-Kyriazi, and Dan Jurafsky. 2022b. [Prompt-and-rerank: A method for zero-shot and few-shot arbitrary textual style transfer with small language models](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 2195–2222.
- Giorgos Vernikos and Andrei Popescu-Belis. 2024. [Don’t rank, combine! combining machine translation hypotheses using quality estimation](#).
- Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc V Le, Ed H. Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. 2023. [Self-consistency improves chain of thought reasoning in language models](#). In *International Conference on Learning Representations*.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, brian ichter, Fei Xia, Ed H. Chi, Quoc V Le, and Denny Zhou. 2022. [Chain of thought prompting elicits reasoning in large language models](#). In *Advances in Neural Information Processing Systems*.
- Sean Welleck, Ximing Lu, Peter West, Faeze Brahman, Tianxiao Shen, Daniel Khashabi, and Yejin Choi. 2023. [Generating sequences by learning to self-correct](#). In *The Eleventh International Conference on Learning Representations*.
- Canwen Xu, Yichong Xu, Shuohang Wang, Yang Liu, Chenguang Zhu, and Julian McAuley. 2023. [Small models are valuable plug-ins for large language models](#).
- 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.
- Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik R Narasimhan, and Yuan Cao. 2023. [React: Synergizing reasoning and acting in language models](#). In *The Eleventh International Conference on Learning Representations*.
- Michihiro Yasunaga, Jure Leskovec, and Percy Liang. 2021. [LM-critic: Language models for unsupervised grammatical error correction](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 7752–7763.
- Chrysoula Zerva, Frédéric Blain, Ricardo Rei, Piyawat Lertvittayakumjorn, José G. C. de Souza, Steffen Eger, Diptesh Kanojia, Duarte Alves, Constantin Orăsan, Marina Fomicheva, André F. T. Martins, and Lucia Specia. 2022. [Findings of the WMT 2022 shared task on quality estimation](#). In *Proceedings of the Seventh Conference on Machine Translation (WMT)*, pages 69–99.
- Biao Zhang, Barry Haddow, and Alexandra Birch. 2023a. [Prompting large language model for machine translation: A case study](#). *arXiv preprint arXiv:2301.07069*.
- Jingqing Zhang, Yao Zhao, Mohammad Saleh, and Peter J. Liu. 2020. [Pegasus: Pre-training with extracted gap-sentences for abstractive summarization](#). In *Proceedings of the 37th International Conference on Machine Learning, ICML’20*.
- Renrui Zhang, Jiaming Han, Aojun Zhou, Xiangfei Hu, Shilin Yan, Pan Lu, Hongsheng Li, Peng Gao, and Yu Jiao Qiao. 2023b. [Llama-adapter: Efficient fine-tuning of language models with zero-init attention](#). *ArXiv*, abs/2303.16199.
- Denny Zhou, Nathanael Schärli, Le Hou, Jason Wei, Nathan Scales, Xuezhi Wang, Dale Schuurmans, Claire Cui, Olivier Bousquet, Quoc V Le, and Ed H.

Chi. 2023. [Least-to-most prompting enables complex reasoning in large language models](#). In *The Eleventh International Conference on Learning Representations*.

A Prompts

We present the few-shot prompts used for all tasks.

Rewrite the input sentence so that it is grammatically accurate.

Source: Yesterday as I was arriving home I saw him in your yard Mack, he looked very satisfied, Curly was laying on the grass and ... "" Yes, Oh Lord he did it because we won his team at football "Nick said."

Target: Yesterday, as I was arriving home, I saw him in your yard, Mack. He looked very satisfied, Curly was lying on the grass and ... "" Yes, oh Lord he did it because we beat his team at football ", Nick said."

Source: According to UNESCO literacy is at the heart of basic education for all and that creating literate environments and societies is essential for achieving the goals of eradicating poverty, reducing child mortality, curbing population growth, achieving gender equality and achieve sustainable development, peace and democracy.

Target: According to UNESCO, literacy is at the heart of basic education for all and creating literate environments and societies is essential for achieving the goals of eradicating poverty, reducing child mortality, curbing population growth, achieving gender equality and achieving sustainable development, peace and democracy.

Source: As you can see I had a very disappointing evening - the worst one during my week's holiday in London.

Target: As you can see I had a very disappointing evening - the worst one of my week's holiday in London.

Source: Anyway, to tell you the truth I'd rather take a train, for instance, it means travelling in a relaxing way, not running risks of accidents, having the chance to read or play "travelling" chess, meeting new people, as in a stage coach but moving faster

Target: Anyway, to tell you the truth, I'd rather take a train. For instance, it means travelling in a relaxing way, not running risks of accidents, having the chance to read or play "travel" chess, meeting new people, like on a stage coach but moving faster.

Source: With the advertisement, you mentioned Mr. Danny Brook and Ms. Tina Truelove were going to play but actually different people were playing, whom I have never seen before.

Target: In the advertisement, you mentioned Mr. Danny Brook and Ms. Tina Truelove were going to perform but actually different people were performing, whom I had never seen before.

Figure 5: LLM prompt for GEC.

Convert the set of key-value attribute pairs in the restaurant domain to a simple English-language text.

Source: name[The Eagle], eatType[coffee shop], food[English], priceRange[high], customer rating[average], area[riverside], familyFriendly[no], near[Burger King]

Target: The Eagle is near Burger King in riverside. It serves expensive English food in a coffee shop setting. It's not child friendly, but has average ratings.

Source: name[Clowns], eatType[coffee shop], food[English], customer rating[5 out of 5], area[riverside], near[Clare Hall]

Target: Clowns is a coffee shop that serves English food and is near Clare Hall. It is located riverside and has a 5 out of 5 customer rating.

Source: name[The Golden Palace], priceRange[more than ÂŁ30], customer rating[high], area[city centre]

Target: The Golden Palace has a high customer rating, with meals costing more than ÂŁ30. It is located in the city center.

Source: name[Wildwood], eatType[coffee shop], food[English], customer rating[1 out of 5], near[Ranch]

Target: Wildwood, English coffee shop, is situated near Ranch and has moderate pricing. It received 1 out of 5 star rating.

Source: name[Taste of Cambridge], eatType[coffee shop], food[English], area[city center], familyFriendly[yes], near[Crowne Plaza Hotel]

Target: Taste of Cambridge is a family-friendly coffee shop that serves English cuisine. It is located in the city center near Crowne Plaza Hotel.

Figure 6: LLM prompt for E2E.

Write a short summary of the article in one sentence.

Source: Many more are feared trapped under rubble after hundreds of buildings collapsed. Thousands of people have been forced to take refuge in temporary shelters and mosques. Some have been left homeless after their houses were destroyed, others have fled their homes amid fears of aftershocks and a possible tsunami. Rescue workers used diggers to remove rubble in their search efforts overnight on Wednesday. Others used their bare hands and shovels to find people. A one survivors were pulled out alive on Wednesday. More than 200 buildings were either seriously damaged or toppled in the earthquake. The Pidie Jaya region, on the north Aceh coast, was the hardest hit. The tremor hit just offshore early on Wednesday morning. Many of the homes in the area have corrugated tin roofs which collapsed. Hundreds have also been rushed to the sole functioning hospital, which has been overwhelmed by patients. Banda Aceh, the provincial capital, was one of the worst hit areas by the 2004 tsunami, caused by acaused by a massive earthquake.
Target: A 6.5-magnitude earthquake struck Aceh province in Indonesia on Wednesday, killing at least 97 people.

Figure 7: LLM prompt for XSum.

The disease has killed nearly 50 people and infected more than 1,400 in Tunisia. = Die Krankheit hat beinahe 50 Menschen getötet und mehr als 1400 Menschen in Tunesien infiziert.

Landray said this failure was particularly exasperating when it came to the use of convalescent plasma, which many doctors believe could have a key role to play in treating seriously ill Covid-19 patients. = Landray beklagt dieses Versagen besonders, wenn es um die Verwendung von rekonvaleszentem Plasma geht, dem laut Meinung vieler Mediziner eine wichtige Rolle bei der Behandlung ernsthaft kranker Covid-19-Patienten zukomme.

Daily cases that numbered in the hundreds dropped to low double digits. = Die täglichen Fälle, die sich auf hunderte beliefen, sanken auf zweistellige Zahlen ab.

However, a recent poll put West at two percent nationwide, neck and neck with the Libertarian Party's Jo Jorgensen and a point ahead of the Green Party's Howie Hawkins. = Jedoch lag West bei einer kürzlich erfolgten Befragung landesweit bei zwei Prozent, Kopf an Kopf mit Jo Jorgensen von der Libertarian Party und einen Punkt vor Howie Hawkins von der Green Party

Scotland's festival scene and sporting events such as the Highland games have been among those affected by restrictions brought in to prevent the spread of Covid-19 = Schottlands Festival- und Sporteventszene, wie die Highland Games, waren unter jenen die von den Einschränkungen, welche eingeführt worden sind um die Ausbreitung von Covid-19 zu verhindern, betroffen waren.

Figure 8: LLM prompt for MT.