

# A Toolbox for Improving Evolutionary Prompt Search

Daniel Griebhaber<sup>a,1,\*</sup>, Maximilian Kimmich<sup>b,1,\*</sup>, Johannes Maucher<sup>a</sup> and Ngoc Thang Vu<sup>b</sup>

<sup>a</sup>Institute for Applied Artificial Intelligence (IAAI), Stuttgart Media University

<sup>b</sup>Institute for Natural Language Processing (IMS), University of Stuttgart

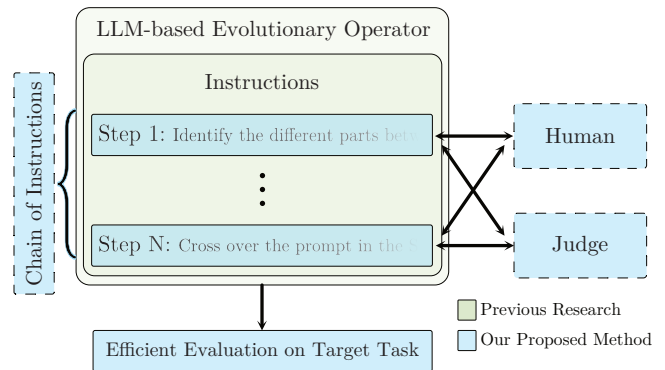
**Abstract.** Evolutionary prompt optimization has demonstrated effectiveness in refining prompts for LLMs. However, existing approaches lack robust operators and efficient evaluation mechanisms. In this work, we propose several key improvements to evolutionary prompt optimization that can partially generalize to prompt optimization in general: 1) decomposing evolution into distinct steps to enhance the evolution and its control, 2) introducing an LLM-based judge to verify the evolutions, 3) integrating human feedback to refine the evolutionary operator, and 4) developing more efficient evaluation strategies that maintain performance while reducing computational overhead. Our approach improves both optimization quality and efficiency. We release our code, enabling prompt optimization on new tasks and facilitating further research in this area.

## 1 Introduction

The recent advent of large language models (LLMs) has ushered in a new era of interactional artificial intelligence, democratizing access to powerful conversational agents and machine translation systems. Despite impressive empirical gains, state-of-the-art LLMs continue to exhibit notable gaps in genuine language understanding, often producing outputs that lack grounding.

One crucial bottleneck in harnessing LLMs for real-world tasks lies in the formulation of *prompts*: the textual instructions that guide model behavior. Prompt quality has been shown to exert a profound influence on model performance, yet designing optimal prompts remains an art: manual tuning is labor-intensive, ad hoc, and often fails to generalize across tasks or domains [3, 16]. To overcome these limitations, a growing body of work has explored automatic prompt optimization methods, including evolutionary strategies in which candidate prompts are iteratively mutated and selected based on LLM responses [28, 5, 27, 7]. While promising, these approaches suffer from two key drawbacks: (i) their feedback loop relies on expensive API calls or compute resources to evaluate every candidate prompt, and (ii) their mutation operators themselves are typically hand-crafted, limiting adaptability and often propagating hallucinations or other undesired artifacts [10, 23].

One promising approach for improving prompt optimization is incorporating human feedback to verify and refine LLM outputs [21]. While LLMs can automate prompt generation, human input remains crucial for verifying the accuracy of their results and correcting errors. By integrating human feedback into the prompt optimization process, we aim to create a more robust system where humans not only verify



**Figure 1:** An overview of the individual components ascribed to our proposed method (blue) compared to only using a one-step instruction for the operator (green).

LLM output but also guide future prompt evolutions. In the case that human feedback is not available (e.g., if there are no domain experts available, or it would be too costly), we allow another LLM to act as judge and to take the responsibility of the human verifying the output of the evolutionary operator. Additionally, since fewer instructions are simpler to verify and chain-of-thought (CoT) reasoning was shown to improve LLM performance as well [31], we believe that the evolutionary operator as well as human feedback and the judging mechanism can benefit from more fine-grained instructions; we call this chain-of-instructions (CoI). This design choice is motivated by several factors: 1) We hypothesize that reducing the complexity of individual instructions makes them easier for the model to follow, akin to how CoT reasoning enhances model performance. 2) It minimizes confusion for the judge, allowing them to assess each step independently, rather than rejecting the entire evolutionary process due to errors in specific parts. 3) Human interaction becomes more efficient, as users can review and validate each evolution step separately – similar to the judge – and provide targeted feedback for improvements. Furthermore, since the evaluation of generated prompts is crucial for the evolution, we introduce methods to efficiently evaluate prompts, finally leading to increased resource efficiency while maintaining task performance.

In this paper, we address the following research questions regarding LLM-based evolutionary prompt optimization methods with the ultimate goal of improving language understanding in LLMs:

RQ1: How can new candidate prompts be evaluated more efficiently without forfeiting overall performance?

RQ2: Does CoI-based evolution help judges and improve the evolutionary operator?

\* Corresponding Authors. E-mails: griebhaber@hdm-stuttgart.de, maximilian.kimmich@ims.uni-stuttgart.de.

<sup>1</sup> Equal contribution.

RQ3: Can LLM-based judges assess the output quality of text-generative evolutionary operators?

RQ4: How can human feedback be effectively leveraged to enhance evolutionary prompt optimization?

RQ5: How does the selection of the LLM influence the effectiveness of the proposed improvements?

To explore these questions, we propose an approach where the LLM serves as both an operator for evolutions and a judge of prompts while allowing humans to intervene when the model output is incorrect. This feedback loop allows human corrections to be leveraged for further optimizing the prompt search process. Additionally, several evaluation strategies aim to make the evaluation more efficient and feedback loops faster. An overview is given in figure 1.

Our contributions are as follows: 1) We introduce a novel human-in-the-loop approach for refining LLM-based prompts, where the feedback is used for future optimizations. 2) We introduce an LLM-based judge for verifying LLM-based evolutions. 3) We present chain-of-instructions as a mechanism to enhance control over the evolutionary process and facilitate more effective feedback. 4) We show how the efficiency of prompt optimization methods can be increased without sacrificing performance. 5) We present empirical results investigating the benefits of our approaches in the context of evolutionary prompt optimization methods. 6) We publish all code that is necessary to reproduce the findings presented in this paper under a permissive license to enable further research into the topic of evolutionary prompt optimization and human feedback.

## 2 Background

In this section, we introduce the foundational concepts and methods that underpin our research on evolutionary prompt optimization.

### 2.1 Large Language Models

LLMs are a class of language models that leverage deep learning techniques to process and generate human language. They have become increasingly popular recently due to their ability to generate high-quality text across a wide range of tasks [3]. LLMs are trained on large corpora of text data, enabling them to learn complex patterns and relationships within language. They have been applied to a variety of natural language processing (NLP) tasks, including text generation, translation, summarization, and question answering.

The input for these LLMs, known as prompts, remains crucial since the LLMs’ performance depends on these prompts. The concept of CoT reasoning has emerged as a powerful approach to enhancing the reasoning capabilities of LLMs. CoT was first introduced by Wei et al. [31], demonstrating that prompting LLMs to decompose complex problems into intermediate reasoning steps significantly improves performance on several tasks.

In scientific and technical domains, the use of LLMs as judges has recently gained increasing attention. Several studies have explored the potential of LLMs for assessing the quality of text in various contexts. For instance, [38] investigated the ability of LLMs to grade academic writing and found that models like GPT-4 can provide feedback comparable to human reviewers. Similarly, works by Liang et al. [14] and Yu et al. [32] have examined LLMs in the context of automated peer review, highlighting their strengths in identifying clarity issues and methodological flaws.

### 2.2 Prompt Optimization

Prompt optimization is the process of finding the most effective prompt for a given task.

For this work, prompt optimization approaches can be classified into two categories: continuous space optimization and discrete prompt optimization methods. The former treat the prompt as a continuous vector, leveraging the fact that the tokens of the prompt are embedded into a continuous space, and optimize it using gradient-based techniques such as gradient descent [26, 17, 16, 33]. However, this approach requires the model to be differentiable with respect to the prompt, which may not always be the case. Furthermore, the resulting prompt lacks interpretability by humans, as the resulting prompt is a continuous vector that cannot easily be mapped to discrete tokens, making it difficult to understand and refine.

Discrete prompt optimization methods, on the other hand, treat the prompt as a discrete sequence of tokens such that the result can easily be observed as a natural language prompt. However, previous methods often rely on gradients over the model parameters to optimize the prompt, which may not be available in black-box models [28, 5, 27, 35]. In contrast, Guo et al. [7] show that prompts can be optimized without gradients, enabling optimization methods for black-box models.

### 2.3 Evolutionary Algorithms

Evolutionary algorithms are a class of optimization techniques inspired by the process of natural selection. They operate by iteratively evolving a population of candidate solutions to a problem, selecting the fittest individuals for reproduction and mutation. Evolutionary algorithms have been successfully applied to a wide range of optimization problems, including function optimization, machine learning, and robotics [8]. Lehman et al. [11] demonstrated that LLMs can be used to perform automatic mutation of prompts, while Meyerson et al. [19] showed the same for crossover operations, paving the way for evolutionary prompt optimization. In the context of prompt optimization, evolutionary algorithms can be used to search for the best prompt for a given task by iteratively evolving and evaluating candidate prompts [4]. Guo et al. [7] is the most related work to ours, as they use a genetic algorithm to optimize prompts for an LLM for NLP tasks.

## 3 Method

Our goal is to find the ideal prompt  $p$  for a given low-resource NLP task (see section 4.1) using in-context learning with an LLM. To achieve this goal, a small set of labeled validation data  $\mathcal{D}$  is available ( $|\mathcal{D}| \leq 200$ ). The LLM is treated as a black-box function, so no access to its parameters or inner architecture is available and needed, enabling models that are only available via APIs.

The optimization process resembles a genetic algorithm, where a population of prompts is evolved over  $T$  generations.

**Initialization:** To generate the initial population, we select the best  $\lfloor I/2 \rfloor$  prompts from a task-specific set of base prompts and generate the remaining  $\lceil I/2 \rceil$  prompts by paraphrasing the selected base prompts, where  $I$  is the population size. This initialization warm-starts the optimization process with a diverse set of prompts [29].

**Evolution:** We follow the work of Lehman et al. [11], Meyerson et al. [19], Guo et al. [7] and use an LLM to perform the operations of *mutation* and *crossover* of the evolutionary algorithm. We use the same operator implementations for Differential Evolution (DE) and Generic Algorithm (GA) as baselines to improve upon, also using

demonstration data in the input for in-context learning. That is, for the classification tasks, we randomly select one input-output pair per class, and a single example for the other tasks.

**Evaluation:** Prompt fitness  $S_i$  is defined as the performance of an evaluation model with prompt  $p_i$  on the validation data:  $S_i = \mathcal{E}(p_i, \mathcal{D})$ .

**Selection:** Prompts for evolution are selected following Lipowski and Lipowska [15] using a roulette-wheel style algorithm with stochastic sampling. Specifically, a prompt is selected with probability  $p_i = \frac{S_i}{\sum_{j=1}^I S_j}$ .

### 3.1 Efficient Evaluation

The optimization process depends on the repeated evaluation of candidate prompts. Although the evaluation set  $\mathcal{D}$  is small, LLM inference is still costly. Given the cost of a single inference step  $c_i$ , the total cost of the evaluation  $c_e$  can be calculated as  $c_e = c_i \times |\mathcal{D}| \times I \times T$ . Since  $c_i$  is fixed to the model, a larger population size  $I$  is generally understood to improve the result of the optimization, and the number of generations  $T$  should be high enough to ensure convergence, we employ and present methods to reduce the overall cost of evaluation without negatively affecting the resulting prompt performance.

#### 3.1.1 Early Stopping

The fitness function normally calculates a mean average score over all samples in  $\mathcal{D}$ . Empirical preliminary examination has shown that the score converges before all samples have been tested. Therefore, we evaluate strategies to reduce the number of evaluation inferences without affecting the resulting prompt performance.

**Moment-based** We propose a moment-based early stopping strategy to stop the evaluation after the score has settled: If the mean absolute difference in evaluation score is less than a minimum change  $\eta_m$  for a window size  $w$ , the evaluation is stopped. The stopping criterion can be expressed with the following inequality:

$$\frac{1}{w} \sum_{n=t-w+1}^t |\mathcal{E}(p_i, \mathcal{D}_n) - \mathcal{E}(p_i, \mathcal{D}_{n-1})| < \eta_m, \\ t \geq w$$

**Parent-based** Except for the first generation, we have access to the performance of all ancestors of  $p_i$  on samples from  $\mathcal{D}$ . We propose to use this information for a parent-based early stopping decision to exit evaluation early if the current prompt  $p_i$  is not performing better by  $\eta_p$  compared to the max score of the parents  $p_a, p_b$  in a sliding window  $w$ . The stopping criterion is then fulfilled if the following inequality is true:

$$\max_{n=t-w+1, \dots, t} (\mathcal{E}(p_i; \mathcal{D}_n) - \max(\mathcal{E}(p_a; \mathcal{D}_n), \mathcal{E}(p_b; \mathcal{D}_n))) < \eta_p, \\ t \geq w$$

We employ the parent-based strategy for generations  $T > 1$  with a fallback to the moment-based strategy when parent performance is not available. Both early stopping methods employ a patience parameter to ignore the first evaluation iterations where the score may change drastically.

#### 3.1.2 Evaluation Strategies

We consider different orderings of  $\mathcal{D}$  for evaluation.

**Shortest First** With the motivation to reduce the number of tokens the LLM needs to process during evaluation, the early stopping strategy can be extended to use an ordered version of  $\mathcal{D}$  that is sorted in ascending order according to the length of the inputs.

**Hardest First** During evolution, the population of prompts is expected to improve. We therefore propose an evaluation strategy that sorts the samples in  $\mathcal{D}$  by the performance of the best parent prompt in ascending order. This is motivated by the fact that a prompt that performs just as well as the best parent prompt on hard samples will not be able to yield a better mean performance on the whole dataset when including samples on which the parent already performed well.

### 3.2 CoI Prompting

We adopt the concept of CoT reasoning for our approach. Rather than instructing the LLM to reason step-by-step, we decompose the instructions for implementing the evolutionary operator into multiple distinct steps. That is, for evolution step  $t$ , we formulate the prompt  $o_t$  to include instructions  $i_t$  and model response  $r_t$  for previous steps as well as the instruction for the current step,  $o_t = i_0, r_0, \dots, i_{t-1}, r_{t-1}, i_t$ . Here, each instruction  $i_t$  is a single operation that the LLM should perform, such as mutating a prompting or crossing over two prompts.

When utilizing demonstration data, we ensure that it aligns with the current stage of evolution, meaning that instructions and model responses up to the current evolution step are included.

### 3.3 Evolution Judge

To avoid an expensive evaluation for prompts that are unlikely to be selected in absence of human feedback, we introduce a judge model  $\mathcal{J}$ . For this, we use another LLM to assess the quality of a prompt candidate  $p_i$  before starting evaluation. To this end, we provide the judge model with the response itself, along with the corresponding inputs that led to it – including demonstration samples, system message and the prompt. In case of CoI, we apply the judge for each evolution step. If the judge model determines a prompt to be of low quality, we ask the evolution model to generate a new response until a predefined number of repetitions is reached. Afterward, if, according to the judge model, there is no prompt of high quality, we continue with a random response from the model.<sup>2</sup>

### 3.4 Human Feedback

To integrate human feedback into the optimization process, we propose a human-in-the-loop approach that actively involves humans at multiple stages of evolutionary prompt optimization. In our framework, human participants are not merely passive evaluators but play an active role in observing, analyzing, and refining the outputs generated by the LLM during each step of the evolutionary process.

Specifically, after each evolutionary step – such as mutation or crossover – humans review the generated model outputs. If deficiencies, ambiguities, or errors are detected, humans intervene by refining the instructions that guide the evolutionary operator. This may include clarifying the language of the instructions, specifying more granular or explicit requirements, or restructuring the sequence of steps

<sup>2</sup> In our implementation, we use the last one since there is no implication on the order of generated model responses, i.e., the randomness is realized via the evolution model.

to reduce confusion for the LLM. Additionally, humans can update or augment the demonstration samples used for in-context learning, ensuring that these examples better illustrate the intended behavior and address previously observed shortcomings.

This process is inherently iterative: after each round of human intervention, the evolutionary process resumes with the updated instructions and demonstration data, allowing for continuous improvement. Over successive cycles, this feedback loop enables the identification and mitigation of persistent weaknesses, such as the LLM’s tendency to overlook subtle distinctions or to generate extraneous output, thereby enhancing the language understanding capabilities of the LLM. By systematically addressing these issues, the overall effectiveness and reliability of the prompt optimization process are enhanced.

Furthermore, this approach allows for the accumulation of best practices and refined instructions, which can be reused or adapted for future tasks or models. The iterative nature of human feedback ensures that the optimization process remains adaptable and responsive to the evolving capabilities and limitations of the underlying LLM. Illustrative examples of how human feedback leads to tangible improvements in prompt optimization are provided in section 6.

## 4 Experimental Setup

This section describes the experiments conducted to explore the effectiveness of the proposed methods. All code, data and information that is necessary to reproduce the results of the presented experiments is published online.<sup>3</sup>

### 4.1 Tasks

We evaluate our proposed method on a wide range of NLP tasks, including sentiment analysis (Stanford Sentiment Treebank (SST) 2 & 5 [30], Movie-Reviews (MR) [18], Customer-Reviews (CR) [9]), subjectivity analysis (Subj [22]), topic classification (AG’s News Topic Classification Dataset (AGNews) [36], Text REtrieval Conference (TREC) [13]), question answering (Stanford Question Answering Dataset (SQuAD) [24]), simplification (A Dataset of Sentence Simplification Evaluation Test (ASSET) [2]) and summarization (Summarizing Arguments in Online Discussions (SAMSum) [6]). For evaluation of prompt fitness during evolution, we select a subset  $\mathcal{D}$  of 200 labeled samples from the validation set. For the final evaluation, we use the whole test set to assess the performance of the evolved prompt.

### 4.2 Evaluation Strategies

To assess the effectiveness of the proposed evaluation strategies, the methods are compared to a baseline where prompt evaluation is performed on the whole validation set  $\mathcal{D}$  and an additional naïve strategy to reduce the evaluation cost by subsampling  $\mathcal{D}$  with a fixed factor. Since the fitness score of a prompt is an important metric in the evolutionary algorithm, we also show the score of the final prompt to preclude negative impacts on the final performance of the optimized prompt.

### 4.3 CoI Prompting

We decompose the instructions of the evolutionary operators, DE and GA, into multiple steps as described in section 3.2. Given the

increased number of instructions required for DE, the CoI-based implementation consists of four steps, whereas GA follows a more concise two-step process. The baseline model for the CoI experiments only performs a single step for the evolution of a prompt. To see the combined effect of CoI and judging, experiments using CoI are performed with and without a judge  $\mathcal{J}$ . For the baseline, the judge can only decide on the single output of the evolution. Experiments are performed on and averaged over all tasks.

### 4.4 Evolution Judge

We employ a judge in our experiments to verify model responses in absence of humans, as described in section 3.3. In detail, the LLM’s instruction is given as follows: "You are acting as a judge. Please read the context, the instruction and the response and decide if the response follows the instruction. If it does, answer 'good'. If it does not, answer 'bad'. Wrap the answer with tags <judgement> and </judgement>. Please also add an explanation for your judgement."<sup>4</sup>. We repeat generating model responses up to three times (if the judge assesses an output as bad).

Similar to the CoI results, the experiments for the judge are compared with and without CoI to assess the combined effect on the final performance.

### 4.5 Human Feedback

After analyzing the evolution model outputs, we iteratively refined the instructions for the evolutionary operator, re-evaluating and further improving the model outputs as needed. Through this process, we applied two consecutive refinements to DE, resulting in DE<sub>1</sub> and DE<sub>2</sub>. Similarly, GA was improved once (GA<sub>1</sub>), as its fewer instructions needed less care. An example of such refinement is provided in section 6.

### 4.6 Hyperparameters

For the evolution, we utilized the quantized version of Llama 3.1 8B Instruct [1] as the generative model. The results are reported after conducting  $T = 10$  generations with a population size of  $I = 10$ . We adopt the approach of Guo et al. [7] for selecting base prompts for the initial population described in section 3. Specifically, depending on the task, we utilize prompts from Mishra et al. [20], Zhang et al. [34], Li et al. [12], Sanh et al. [25], Zhang et al. [37]. However, for SQuAD, we employ a single manually crafted prompt alongside generated prompts obtained using the forward mode generation method proposed by Zhou et al. [39].

For both paraphrasing and evolutionary steps, sampling was applied in decoding using a temperature of  $t = 0.5$  to increase output variance. We used the same model for the judge and for the evaluation, i.e., Llama 3.1 8B Instruct, but with greedy decoding for increased correctness.

For the early stopping, we set  $\eta_m = 10^{-3}$ , the window size  $w = 10$ ,  $\eta_p = 10^{-3}$  and a patience of 20.

<sup>3</sup> <https://gitlab.mi.hdm-stuttgart.de/griesshaber/evoprompt>

<sup>4</sup> This prompt is the result of a manual investigation; since this work deals with the optimization of prompts and the manual tuning of prompts contradicts the motivation of this work, this manual tuning is only necessary at this point, but neither reflects our intention to improve prompt optimization, nor does it counteract it.

**Table 1:** Baseline results for our hyperparameters and models for various classification, question answering, and text generation tasks. Results are reported for two evolutionary algorithms: Differential Evolution (DE) and Genetic Algorithm (GA). GA outperforms DE on most tasks, including AGNews, ASSET, SAMSum, SQuAD, SST-2, SST-5, and Subj. DE achieves better results on CR, MR, and TREC. This indicates that – without any of our improvements – while GA generally yields better overall performance, DE remains competitive on certain tasks.

Task	Evolution Algorithm	
	DE	GA
AGNews	86.89%	<b>87.99%</b>
ASSET	54.74%	<b>56.78%</b>
CR	<b>91.85%</b>	91.40%
MR	<b>91.30%</b>	90.55%
SAMSum	29.55%	<b>29.82%</b>
SQuAD	86.72%	<b>89.27%</b>
SST 2	94.51%	<b>95.50%</b>
SST 5	56.24%	<b>56.33%</b>
Subj	80.75%	<b>84.10%</b>
TREC	<b>83.20%</b>	78.80%

#### 4.7 Other LLMs

To explore the effect the choice of LLM has on the performance of the proposed methods, we conduct experiments with numerous other LLMs :

- LI3.1 (8B): Llama 3.1 8B Instruct<sup>5</sup> (our base model).
- MI (7B): Mistral Instruct 7B<sup>6</sup>
- Q2.5 (7B): Qwen 2.5 7B Instruct<sup>7</sup>
- R1Q (1.5B): Deep Seek R1 Distill Qwen 1.5B Instruct<sup>8</sup>
- R1Q (7B): Deep Seek R1 Distill Qwen 7B Instruct<sup>9</sup>
- GI (1B): Google Gemma 1B Instruct<sup>10</sup>
- GI (7B): Google Gemini 7B Instruct<sup>11</sup>

## 5 Quantitative Results

This section presents the results of the conducted experiments in a quantitative manner to show the effectiveness of each modification, analyzing their implications with respect to our research questions.

### 5.1 Baseline

Since we aim to investigate the impact of our proposed extensions to evolutionary prompt optimization, table 1 presents baseline results for our chosen hyperparameters and LLM models for all presented tasks and both evolution algorithms. The results mostly replicate the final results presented in Guo et al. [7], but with deviations in hyperparameters and implementation.

### 5.2 RQ1: Efficient Evaluation

Table 2 shows the results for the experimental evaluation of the proposed efficient evaluation strategies. The results indicate that, while

Model references on Hugging Face:

<sup>5</sup> meta-llama/Meta-Llama-3.1-8B-Instruct

<sup>6</sup> mistralai/Mistral-7B-Instruct-v0.1

<sup>7</sup> Qwen/Qwen2.5-7B-Instruct

<sup>8</sup> deepseek-ai/DeepSeek-R1-Distill-Qwen-1.5B

<sup>9</sup> deepseek-ai/DeepSeek-R1-Distill-Qwen-7B

<sup>10</sup> google/gemma-1b-it

<sup>11</sup> google/gemma-7b-it

**Table 2:** Comparing the different evaluation strategies to the baseline: average difference of evaluation scores ( $\Delta S$ ), difference in number of tokens used during prompt evaluation ( $\Delta c_e$ ), token count ( $c_e/c_b$ ) and runtime ( $t_e/t_b$ ) of the experiment as a fraction of the baseline. All reported values are averaged over tasks.

Strategy	$\Delta S$	$\Delta c_e$	$c_e/c_b$	$t_e/t_b$
Subsample	-2.28%	-7.3M	33.5%	55.6%
Early Stopping	-1.43%	-7.6M	31.0%	53.2%
Shortest First	<b>+0.11%</b>	-7.9M	28.3%	43.6%
Hardest First	-0.50%	<b>-8.0M</b>	<b>25.5%</b>	<b>42.8%</b>

**Table 3:** Performance improvements of CoI over baselines with ( $\checkmark$ ) and without ( $\times$ ) using a judge assessing model outputs. Incorporating CoI consistently improves results across both Differential Evolution (DE) and Genetic Algorithm (GA), with a judge yielding best results for both algorithms. For DE, the mean improvement increases from +1.20% to +1.73%, with a higher maximum gain. Similarly, GA sees an increase in mean improvement from +0.68% to +1.00%, with a notable maximum gain of +4.40%. These results demonstrate that CoI optimization yields more reliable and higher-quality model outputs.

Evolution Algorithm	$\mathcal{J}$	$\Delta S$		
		mean	min	max
DE	$\times$	+1.20%	-0.35%	+3.00%
	$\checkmark$	<b>+1.73%</b>	-0.80%	+3.56%
GA	$\times$	+0.68%	-0.88%	+2.55%
	$\checkmark$	<b>+1.00%</b>	-0.63%	+4.40%

all evaluation strategies reduce the number of tokens needed to score the candidate prompts, the naïve approach of subsampling  $\mathcal{D}$  performs in average worst, demonstrating the importance of high-quality prompt candidate scores for the evolutionary algorithm. Notably, the suggested strategies of evaluating on the shortest and hardest samples first only show minor deviation from the baseline scores while also being most effective reducing used tokens and runtime, whereas early stopping on the unordered scoring set may decrease performance. While the *Hardest First* strategy reduces the evaluation cost the most, *Shortest First* is the only strategy that did not show any decrease in the final evaluation. In total, to answer RQ1, both strategies, Hardest First and Shortest First, can effectively reduce the compute usage. They should be chosen based on individual preferences on the final task performance.

### 5.3 RQ2&3: CoI Prompting & Evolution Judge

In table 3, the relative improvements that can be achieved with CoI prompting are presented. The average across tasks is positive for all configurations, regardless of evolution algorithm and the use of an additional judge.

Notably, the mean improvement for the DE algorithm is higher than for the GA algorithm, independent of whether the judge is used. Since DE is more complex with a higher number of steps, this indicates that CoI helps by breaking the algorithm into discrete steps that can be performed individually. Furthermore, CoI yields the best performance in combination with the judge, additionally motivating to verify model outputs automatically in LLM-based evolutionary operators.

Similar to the enhancements achieved with CoI, incorporating the judge in our approach consistently outperforms the baseline methods, as demonstrated in table 4, regardless of whether CoI is utilized. This result indicates that the judge can successfully detect and reject prompts which are determined to be of low quality before evaluation and therefore provides a positive answer to RQ3. The resulting increase in the number of high-quality prompts in the population before

**Table 4:** Performance improvements of applying a judge to assess evolution with (✓) and without (✗) using CoI. The judge improves average performance ( $\Delta S$ ) across both Differential Evolution (DE) and Genetic Algorithm (GA). With CoI, DE sees a higher mean improvement (+0.87% vs. +0.34%) and a reduced worst-case drop in performance. GA also benefits, showing a larger average gain (+0.97% vs. +0.65%) and achieving the highest observed improvement overall (+3.20%). These results highlight the effectiveness of using a judge, and especially in conjunction with CoI, to guide the evolutionary process.

Evolution Algorithm	CoI	$\Delta S$		
		mean	min	max
DE	✗	+0.34%	-0.86%	+2.80%
	✓	<b>+0.87%</b>	-0.18%	+2.58%
GA	✗	+0.65%	-0.22%	+1.79%
	✓	<b>+0.97%</b>	-0.35%	+3.20%

selection seems to yield an overall improvement in the performance of the final evolved prompt as observed in our results.


In combination, the results from tables 3 and 4 show – in the ablation cases where either CoI or the judge are removed – that both concepts work best in combination, providing a positive answer to RQ2. This is to be expected, since the decomposition of the evolutionary prompt also allows the judge to assess each smaller step separately, compared to judging the whole output including multiple steps over longer input-output pairs.

**Table 5:** Relative score improvements of evolution strategies revised using human feedback (the subscript indicates the iteration) compared to DE and GA, respectively. Incorporating human feedback yields consistent performance gains across most tasks. On average, the second iteration of DE ( $DE_2$ ) shows the highest mean improvement (+1.67%) reflecting the notion of consecutive refinements, followed by the first DE refinement ( $DE_1$ , +1.11%) and the GA refinement ( $GA_1$ , +0.75%). The largest individual improvements are observed on TREC and ASSET, indicating that human feedback is particularly effective for tasks involving question classification and text simplification.

Task	$\Delta S_{DE_1}$	$\Delta S_{DE_2}$	$\Delta S_{GA_1}$
AGNews	+1.03%	+1.54%	+0.20%
ASSET	+1.74%	+2.36%	-0.18%
CR	+1.55%	+2.10%	+2.15%
MR	+0.15%	+0.15%	-0.05%
SAMSum	-0.19%	+1.09%	+0.43%
SQuAD	+1.09%	+0.95%	+0.84%
SST 2	+0.90%	+1.68%	+0.88%
SST 5	+1.43%	+2.06%	+0.81%
Subj	+0.20%	+1.40%	+0.20%
TREC	+3.20%	+3.40%	+2.20%
Mean	<b>+1.11%</b>	<b>+1.67%</b>	<b>+0.75%</b>

#### 5.4 RQ4: Human Feedback

The results presented in table 5 demonstrate the relative performance improvements of evolution strategies incorporating human feedback ( $DE_1$ ,  $DE_2$ , and  $GA_1$ ) over their respective baseline methods (DE and GA) across multiple tasks. Overall,  $DE_2$  consistently outperforms its baseline, achieving the highest mean improvement of +1.67%, compared to +1.11% for  $DE_1$ .  $GA_1$  also benefits from human feedback, but shows a more modest mean improvement of +0.75%. While most tasks exhibit performance gains, there are a few instances where minimal or negative changes occur. Notably, the largest improvements are observed on the TREC dataset showing substantial gains, with  $DE_2$  achieving the highest relative improvement of +3.40%. However, there



**Figure 2:** Box plot illustrating the quantitative effectiveness of various LLMs based on performance metrics across our evaluation set. The models are listed in section 4.7. The y-axis represents the relative improvement in performance if CoI and the judge are used. Mean performance is consistently improved across all tested models, with Gemma profiting the most.

is a variance on individual tasks induced by the randomness of the evolution, but these findings suggest that integrating human feedback into evolution strategies can enhance performance, with DE-based approaches benefiting more noticeably than GA-based ones.<sup>12</sup>

Following this evaluation and to answer RQ4, we can say that the evolutionary operator is effectively improved using human feedback as proposed in our approach.

#### 5.5 RQ5: Effectiveness of different LLMs

Figure 2 shows the performance difference between a baseline and runs using both the judge and CoI, with different LLMs models. The evaluation was performed for all tasks presented in 4.1. The results indicate that the proposed methods are effective across all tested models with an overall positive average improvement. Interestingly, the smaller models, such as R1Q (1.5B) and GI (1B), show a higher variance in performance, especially when compared directly to the larger variants of the same model (R1Q (7B) and GI (7B) respectively).

To answer RQ5, we were able to demonstrate the effectiveness of our proposed method across a wide combination of tasks and models.

## 6 Qualitative Analysis of Human Feedback

Figure 3 presents an example of human feedback within our approach. Upon reviewing the model output from the first evolution step, it was observed that not all differences were identified, and some similarities were erroneously included. Additionally, unexpected extraneous output was generated. To address these issues, the prompt for the first evolution step was refined by incorporating the following instructions: "Output a list of all different parts and make sure that differences are only in the form of words and phrases." and "If the same phrase appears in both prompts, do not list it, i.e., do not list similarities."

This example also highlights the ambiguities associated with human involvement: multiple instructions exist, just as there are various

<sup>12</sup> We note that the combination of judge-based and human feedback – where humans corrected the model output if it was judged non-compliant with the instructions – did not consistently enhance performance across all tasks. Consequently, we have opted not to present these results.

### Input:

Step 1: Identify the different parts between Prompt 1 and Prompt 2:

Prompt 1: Analyze the sentence and categorize it into one of five categories based on the sentiment: terrible, bad, okay, good, or great.

Prompt 2: Classify the given review into one of five categories: extremely negative (terrible), somewhat negative (bad), neutral (okay), somewhat positive (good), or extremely positive (great).

### Response:

1. Identifying the different parts between Prompt 1 and Prompt 2:

Prompt 1: Analyze the sentence and categorize it into one of five categories based on the sentiment: terrible, bad, okay, good, or great.

Prompt 2: Classify the given review into one of five categories: extremely negative (terrible), somewhat negative (bad), neutral (okay), somewhat positive (good), or extremely positive (great).

Different parts:

- "sentence" vs "review"

- "analyze" vs "classify"

- "terrible, bad, okay, good, or great" vs "extremely negative (terrible), somewhat negative (bad), neutral (okay), somewhat positive (good), or extremely positive (great)"

- "based on the sentiment" only appears in Prompt 1

- ~~"based on the sentiment" vs "based on the sentiment" (same phrase)~~

- ~~"five categories" vs "five categories" (same phrase)~~

- ~~The only difference is the subject of analysis, which is a sentence in Prompt 1 and a review in Prompt 2.~~

**Figure 3:** An example for the first step of evolution for DE: The expected response mentions mutations for all spotted differences (marked in red) and omits the similarities as well as the last statement, which is evidently wrong (marked using strikethrough in red). Demonstration samples for in-context learning are omitted for clarity.

possible prompts for a task. By iteratively analyzing the evolution model’s output and refining the instructions accordingly, we can facilitate human feedback, ultimately enhancing the prompt optimization process.

Finally, we claim that inspecting the model output and adapting the instructions correspondingly can be accomplished in about half an hour, which is a reasonable time investment for the performance improvements achieved in our experiments.

## 7 Conclusion

We introduced and investigated extensions to evolutionary prompt optimization that leverage CoI, an LLM-based judge, human feedback and efficient evaluation methods to optimize prompts for a given task.

CoI, by enabling greater control and better decision-making, can enhance performance in prompt optimization (RQ2) and holds promise for broader applications. In particular, when combined with judge-based assessment (RQ3) and human feedback (RQ4), it provides a robust framework for identifying optimal prompts in NLP tasks. Lastly, beyond performance improvements, reducing computational cost is also a key consideration. Our efficient evaluation methods offer a significant reduction in computational overhead while maintaining performance during the search for optimal prompts (RQ1).

We are convinced that our contributions, including investigations and releasing our code, help future research in the area, promoting the effective and efficient use of LLMs in NLP, and especially help in grounding LLMs for better language understanding.

## 8 Limitations

In this work we only focus on the optimization of the prompts while not focusing on optimizing the verbalizer extracting the predictions for the tasks, which could be a potential improvement since the prompt can contain directives as to what to expect in the model output.

Furthermore, running multiple experiments on the same task can lead to different results due to the stochastic nature of LLMs and the

evolutionary algorithms, providing a more reliable performance estimate. However, since the experiments are time-intensive, we instead mitigate this effect by averaging the results over multiple tasks. This also allows us to analyze the performance of our methods across a wide range of tasks, but may not be representative for individual tasks.

Also, although we optimize for faster runtimes and lower token usage, LLMs still require large amounts of compute resources and energy which potentially makes the methods and results presented in this paper inaccessible to some groups without access to such resources. For example, a single optimization of a prompt for SAMSum using the *hardest first* strategy needed 4:24h on a single NVIDIA A6000 GPU while the average GPU memory consumption was about 20GB.

## References

- [1] AI@Meta. The Llama 3 herd of models, 2024. URL <https://arxiv.org/abs/2407.21783>.
- [2] F. Alva-Manchego, L. Martin, A. Bordes, C. Scarton, B. Sagot, and L. Specia. ASSET: A dataset for tuning and evaluation of sentence simplification models with multiple rewriting transformations. In D. Jurafsky, J. Chai, N. Schluter, and J. Tetreault, editors, *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4668–4679, Online, jul 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.acl-main.424.
- [3] T. Brown, B. Mann, N. Ryder, M. Subbiah, J. D. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, S. Agarwal, A. Herbert-Voss, G. Krueger, T. Henighan, R. Child, A. Ramesh, D. Ziegler, J. Wu, C. Winter, C. Hesse, M. Chen, E. Sigler, M. Litwin, S. Gray, B. Chess, J. Clark, C. Berner, S. McCandlish, A. Radford, I. Sutskever, and D. Amodei. Language models are few-shot learners. In H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, pages 1877–1901. Curran Associates, Inc., 2020. URL <https://dl.acm.org/doi/abs/10.5555/3495724.3495883>.
- [4] A. Chen, D. M. Dohan, and D. R. So. EvoPrompting: language models for code-level neural architecture search. In *Proceedings of the 37th International Conference on Neural Information Processing Systems*, NIPS ’23, Red Hook, NY, USA, 2023. Curran Associates Inc. URL <https://dl.acm.org/doi/10.5555/3666122.3666464>.
- [5] M. Deng, J. Wang, C.-P. Hsieh, Y. Wang, H. Guo, T. Shu, M. Song, E. Xing, and Z. Hu. RLPrompt: Optimizing discrete text prompts with

- reinforcement learning. In Y. Goldberg, Z. Kozareva, and Y. Zhang, editors, *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 3369–3391, Abu Dhabi, United Arab Emirates, dec 2022. Association for Computational Linguistics. doi: 10.18653/v1/2022.emnlp-main.222.
- [6] B. Gliwa, I. Mochol, M. Biesek, and A. Wawer. SAMSum corpus: A human-annotated dialogue dataset for abstractive summarization. In L. Wang, J. C. K. Cheung, G. Carenini, and F. Liu, editors, *Proceedings of the 2nd Workshop on New Frontiers in Summarization*, pages 70–79, Hong Kong, China, nov 2019. Association for Computational Linguistics. doi: 10.18653/v1/D19-5409.
- [7] Q. Guo, R. Wang, J. Guo, B. Li, K. Song, X. Tan, G. Liu, J. Bian, and Y. Yang. Connecting large language models with evolutionary algorithms yields powerful prompt optimizers. In *The Twelfth International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=ZG3RaNI8O8>.
- [8] J. H. Holland. *Adaptation in Natural and Artificial Systems: An Introductory Analysis with Applications to Biology, Control and Artificial Intelligence*. MIT Press, Cambridge, MA, USA, 1992. ISBN 0262082136.
- [9] M. Hu and B. Liu. Mining and summarizing customer reviews. In *Proceedings of the Tenth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '04*, page 168–177, New York, NY, USA, 2004. Association for Computing Machinery. ISBN 1581138881. doi: 10.1145/1014052.1014073.
- [10] Z. Ji, N. Lee, R. Frieske, T. Yu, D. Su, Y. Xu, E. Ishii, Y. J. Bang, A. Madotto, and P. Fung. Survey of hallucination in natural language generation. (12), mar 2023. ISSN 0360-0300. doi: 10.1145/3571730.
- [11] J. Lehman, J. Gordon, S. Jain, K. Ndousse, C. Yeh, and K. O. Stanley. Evolution through large models, 2022. URL <https://arxiv.org/abs/2206.08896>.
- [12] B. Li, R. Wang, J. Guo, K. Song, X. Tan, H. Hassan, A. Menezes, T. Xiao, J. Bian, and J. Zhu. Deliberate then generate: Enhanced prompting framework for text generation. 2023.
- [13] X. Li and D. Roth. Learning question classifiers. In *COLING 2002: The 19th International Conference on Computational Linguistics*, 2002. URL <https://aclanthology.org/C02-1150/>.
- [14] W. Liang, Y. Zhang, H. Cao, B. Wang, D. Ding, X. Yang, K. Vodrahalli, S. He, D. S. Smith, Y. Yin, D. A. McFarland, and J. Zou. Can large language models provide useful feedback on research papers? a large-scale empirical analysis. 2023.
- [15] A. Lipowski and D. Lipowska. Roulette-wheel selection via stochastic acceptance. (6):2193–2196, 2012. ISSN 0378-4371. doi: 10.1016/j.physa.2011.12.004.
- [16] P. Liu, W. Yuan, J. Fu, Z. Jiang, H. Hayashi, and G. Neubig. Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing. (9), jan 2023. ISSN 0360-0300. doi: 10.1145/3560815.
- [17] X. Liu, K. Ji, Y. Fu, W. Tam, Z. Du, Z. Yang, and J. Tang. P-tuning: Prompt tuning can be comparable to fine-tuning across scales and tasks. In S. Muresan, P. Nakov, and A. Villavicencio, editors, *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 61–68, Dublin, Ireland, may 2022. Association for Computational Linguistics. doi: 10.18653/v1/2022.acl-short.8.
- [18] A. L. Maas, R. E. Daly, P. T. Pham, D. Huang, A. Y. Ng, and C. Potts. Learning word vectors for sentiment analysis. In D. Lin, Y. Matsumoto, and R. Mihalcea, editors, *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 142–150, Portland, Oregon, USA, jun 2011. Association for Computational Linguistics. URL <https://aclanthology.org/P11-1015/>.
- [19] E. Meyerson, M. J. Nelson, H. Bradley, A. Gaier, A. Moradi, A. K. Hoover, and J. Lehman. Language model crossover: Variation through few-shot prompting. (4), 2024. doi: 10.1145/3694791.
- [20] S. Mishra, D. Khashabi, C. Baral, and H. Hajishirzi. Cross-task generalization via natural language crowdsourcing instructions. In *Annual Meeting of the Association for Computational Linguistics*, 2021.
- [21] L. Ouyang, J. Wu, X. Jiang, D. Almeida, C. L. Wainwright, P. Mishkin, C. Zhang, S. Agarwal, K. Slama, A. Ray, J. Schulman, J. Hilton, F. Kelton, L. Miller, M. Simens, A. Askell, P. Welinder, P. Christiano, J. Leike, and R. Lowe. Training language models to follow instructions with human feedback. In *Proceedings of the 36th International Conference on Neural Information Processing Systems, NIPS '22*, Red Hook, NY, USA, 2022. Curran Associates Inc. ISBN 9781713871088.
- [22] B. Pang, L. Lee, and S. Vaithyanathan. Thumbs up? sentiment classification using machine learning techniques. In *Proceedings of the 2002 Conference on Empirical Methods in Natural Language Processing (EMNLP 2002)*, pages 79–86. Association for Computational Linguistics, jul 2002. doi: 10.3115/1118693.1118704.
- [23] G. Perković, A. Drobnjak, and I. Botički. Hallucinations in LLMs: Understanding and addressing challenges. In *2024 47th MIPRO ICT and Electronics Convention (MIPRO)*, pages 2084–2088, 2024. doi: 10.1109/MIPRO60963.2024.10569238.
- [24] P. Rajpurkar, J. Zhang, K. Lopyrev, and P. Liang. SQuAD: 100,000+ questions for machine comprehension of text. In J. Su, K. Duh, and X. Carreras, editors, *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2383–2392, Austin, Texas, nov 2016. Association for Computational Linguistics. doi: 10.18653/v1/D16-1264.
- [25] V. Sanh, A. Webson, C. Raffel, S. H. Bach, L. Sutawika, Z. Alyafeai, A. Chaffin, A. Stiegler, T. L. Scao, A. Raja, M. Dey, M. S. Bari, C. Xu, U. Thakker, S. Sharma, E. Szczechla, T. Kim, G. Chhablani, N. V. Nayak, D. Datta, J. D. Chang, M. T.-J. Jiang, H. Wang, M. Manica, S. Shen, Z.-X. Yong, H. Pandey, R. Bawden, T. Wang, T. Neeraj, J. Rozen, A. Sharma, A. Santilli, T. Févry, J. A. Fries, R. Teehan, S. Biderman, L. Gao, T. Bers, T. Wolf, and A. M. Rush. Multitask prompted training enables zero-shot task generalization. 2021.
- [26] T. Schick and H. Schütze. Exploiting Cloze-questions for few-shot text classification and natural language inference. In P. Merlo, J. Tiedemann, and R. Tsarfay, editors, *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 255–269, Online, apr 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.eacl-main.20.
- [27] W. Shi, X. Han, H. Gonen, A. Holtzman, Y. Tsvetkov, and L. Zettlemoyer. Toward human readable prompt tuning: Kubrick’s the shining is a good movie, and a good prompt too? In H. Bouamor, J. Pino, and K. Bali, editors, *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 10994–11005, Singapore, dec 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.findings-emnlp.733.
- [28] T. Shin, Y. Razeghi, R. L. Logan IV, E. Wallace, and S. Singh. AutoPrompt: Eliciting Knowledge from Language Models with Automatically Generated Prompts. In B. Webber, T. Cohn, Y. He, and Y. Liu, editors, *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4222–4235, Online, nov 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.emnlp-main.346.
- [29] D. So, Q. Le, and C. Liang. The evolved transformer. In K. Chaudhuri and R. Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning, Proceedings of Machine Learning Research*, pages 5877–5886. PMLR, 09–15 Jun 2019. URL <https://proceedings.mlr.press/v97/so19a.html>.
- [30] R. Socher, A. Perelygin, J. Wu, J. Chuang, C. D. Manning, A. Ng, and C. Potts. Recursive deep models for semantic compositionality over a sentiment treebank. In D. Yarowsky, T. Baldwin, A. Korhonen, K. Livescu, and S. Bethard, editors, *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1631–1642, Seattle, Washington, USA, oct 2013. Association for Computational Linguistics. URL <https://aclanthology.org/D13-1170/>.
- [31] J. Wei, X. Wang, D. Schuurmans, M. Bosma, B. Ichter, F. Xia, E. H. Chi, Q. V. Le, and D. Zhou. Chain-of-thought prompting elicits reasoning in large language models. In *Proceedings of the 36th International Conference on Neural Information Processing Systems, NIPS '22*, Red Hook, NY, USA, 2022. Curran Associates Inc. ISBN 9781713871088. URL <https://dl.acm.org/doi/10.5555/3600270.3602070>.
- [32] J. Yu, Z. Ding, J. Tan, K. Luo, Z. Weng, C. Gong, L. Zeng, R. Cui, C. Han, Q. Sun, Z. Wu, Y. Lan, and X. Li. Automated peer reviewing in paper SEA: Standardization, evaluation, and analysis. In *Conference on Empirical Methods in Natural Language Processing*, 2024.
- [33] N. Zhang, L. Li, X. Chen, S. Deng, Z. Bi, C. Tan, F. Huang, and H. Chen. Differentiable prompt makes pre-trained language models better few-shot learners. In *International Conference on Learning Representations*, 2022. URL <https://openreview.net/forum?id=ek9a0qlafW>.
- [34] S. Zhang, S. Roller, N. Goyal, M. Artetxe, M. Chen, S. Chen, C. Dewan, M. T. Diab, X. Li, X. V. Lin, T. Mihaylov, M. Ott, S. Shleifer, K. Shuster, D. Simig, P. S. Koura, A. Sridhar, T. Wang, and L. Zettlemoyer. OPT: Open pre-trained transformer language models. 2022.
- [35] T. Zhang, X. Wang, D. Zhou, D. Schuurmans, and J. E. Gonzalez. TEMPERA: Test-time prompt editing via reinforcement learning. In *The Eleventh International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=gSHyqBijPFO>.
- [36] X. Zhang, J. Zhao, and Y. LeCun. Character-level convolutional networks for text classification. In *Proceedings of the 29th International Conference on Neural Information Processing Systems - Volume 1, NIPS '15*, page 649–657, Cambridge, MA, USA, 2015. MIT Press.
- [37] Y. Zhang, L. Cui, D. Cai, X. Huang, T. Fang, and W. Bi. Multi-task instruction tuning of llama for specific scenarios: A preliminary study on writing assistance. 2023.



- [38] L. Zheng, W.-L. Chiang, Y. Sheng, S. Zhuang, Z. Wu, Y. Zhuang, Z. Lin, Z. Li, D. Li, E. P. Xing, H. Zhang, J. E. Gonzalez, and I. Stoica. Judging LLM-as-a-judge with MT-Bench and chatbot arena. 2023.
- [39] Y. Zhou, A. I. Muresanu, Z. Han, K. Paster, S. Pitis, H. Chan, and J. Ba. Large language models are human-level prompt engineers. 2022.