

# CriSPO: Multi-Aspect Critique-Suggestion-guided Automatic Prompt Optimization for Text Generation

Han He\*, Qianchu Liu\*, Lei Xu\*, Chaitanya Shivade,  
Yi Zhang, Sundararajan Srinivasan, Katrin Kirchhoff

Amazon AWS AI Labs  
{hankcs, liufqian, leixx, shivadc,  
yizhngn, sundarsr, katrinki}@amazon.com

## Abstract

Existing automatic prompt engineering methods are typically designed for discriminative tasks, where new task prompts are iteratively refined with limited feedback from a single metric reflecting a single aspect. However, these approaches are suboptimal for generative tasks, which require more nuanced guidance beyond a single numeric metric to improve the prompt and optimize multiple aspects of the generated text. To address these challenges, we propose a novel multi-aspect Critique-Suggestion-guided automatic Prompt Optimization (CriSPO) approach. CriSPO introduces a *critique-suggestion module* as its core component. This module spontaneously discovers aspects, and compares generated and reference texts across these aspects, providing specific suggestions for prompt modification. These clear critiques and actionable suggestions guide a *receptive optimizer module* to make more substantial changes, exploring a broader and more effective search space. To further improve CriSPO with multi-metric optimization, we introduce an Automatic Suffix Tuning (AST) extension to enhance the performance of task prompts across multiple metrics. We evaluate CriSPO on 4 state-of-the-art Large Language Models (LLMs) across 4 summarization and 5 Question Answering (QA) datasets. Extensive experiments show 3-4% ROUGE score improvement on summarization and substantial improvement of various metrics on QA.

**Code** — <https://github.com/amazon-science/CriSPO>

**Extended version** — <https://arxiv.org/abs/2410.02748>

## 1 Introduction

LLMs have emerged as powerful tools for various natural language processing tasks, including text generation (Brown et al. 2020). To fully leverage their capabilities, a critical step is to design a precise *task prompt* which specifies the desired behavior of the LLM to solve a task. Manual prompt engineering is often laborious, skill-intensive and sub-optimal, motivating the need for automatic prompt engineering techniques which automatically tune the task prompt.

Recent research has made notable progress in automatic prompt engineering for discriminative tasks, such as text classification (Zhou et al. 2022; Yang et al. 2023; Pryzant

et al. 2023; Sordoni et al. 2024). These methods focus on optimizing task prompts for a single metric on a single aspect. The process typically involves instructing an LLM optimizer with a meta-prompt to generate new task prompts based on previously sampled task prompts and their corresponding scores. By iteratively exploring candidates and selecting the task prompt with the highest score, performance on the target metric improves over numerous iterations. However, applying these methods directly to text generation tasks, such as summarization, is sub-optimal due to challenges in obtaining effective *optimization signals*. Unlike classification tasks, where metrics are straightforward (eg. accuracy), automatic metrics for text generation, like ROUGE (Lin 2004), provides limited guidance for prompt refinement. For example, a lower ROUGE score may result from aspects such as mismatched length, differences in word choice due to formality, or varying writing formats, making it difficult to guide LLMs in prompt modification without fine-grained feedback targeting these individual aspects. Furthermore, evaluating text generation involves multiple metrics (Fabbri et al. 2021; Gao and Wan 2022; Elangovan et al. 2024). In addition to reference similarity, other metrics such as factual consistency, which can be assessed using metrics like AlignScore (Zha et al. 2023), are also important. Balancing or utilizing these multiple metrics is not fully addressed by existing prompt engineering methods that focus on optimizing a single metric.

To address these challenges, we introduce CriSPO, a multi-aspect Critique-Suggestion-guided automatic Prompt Optimization (CriSPO) approach. Overall, our approach employs LLMs to automatically identifies multi-aspect prompt revision suggestions, based on which prompts are automatically designed and refined (Table 8 in Appendix shows a working example of how a prompt gets revised in CriSPO). Inspired by recent self-reflection studies, where LLMs generate verbal feedback to aid in self-improvement (Gero et al. 2023; Shinn et al. 2023; Madaan et al. 2024), we designed the first key component of CriSPO: the **multi-aspect critique-suggestion meta-prompt**. It automatically discovers proper aspects to compare generated text with reference, write critiques of flaws (Pryzant et al. 2023) and suggestions to improve the task prompt (Figure 2 shows a word cloud of aspects identified by CriSPO, including number of words, style, and precision). Both critiques and suggestions, written in natural language, are more helpful for prompt improvement

\*These authors contributed equally.

than a single ROUGE score. We then create a **receptive optimizer meta-prompt** that generates new prompts. In addition to conditioning on previous high-score task prompts and scores, this optimizer also reviews the past critiques and suggestions. It then generates an overall suggestion and an improved task prompt candidate in a Chain-of-Thought (CoT) (Wei et al. 2022) manner. Our approach iteratively optimizes the task prompt using LLMs similar to previous work like Optimization by PROMpting (OPRO) (Yang et al. 2023), but it enriches the training signal with multi-aspect critiques and suggestions to better optimize a text generation metric. To further enhance performance by allowing the prompt to access external data, we design the **task prompt template** that contains placeholders for In-Context Learning (ICL) examples or retrieved contexts. The receptive optimizer meta-prompt generates these templates directly, so it can flexibly move components in task prompt for better organization.

While CriSPO offers multi-aspect guidance for optimizing text generation through critiques and suggestions, we further enhance this guidance by incorporating multiple metrics as additional teaching signals. To this end, we propose a novel Automatic Suffix Tuning (AST) extension which divides prompts into chunks conquering different metrics. Through multi-objective learning, we improve each new metric with little to no drop in existing metrics.

We test CriSPO on state-of-the-art LLMs, including Claude (Anthropic 2023, 2024), Mistral (Jiang et al. 2023) and Llama3 (MetaAI 2024), across 9 heterogeneous datasets. These include 4 summarization datasets spanning various abstractiveness, formats, and domains, as well as 5 QA datasets. Extensive experiments demonstrate that CriSPO significantly improves prompt quality and task performance over strong baselines as verified by human evaluation. We also conduct ablation study to assess the effectiveness of key ingredients.

**Our contributions are summarized below:**

1) We propose CriSPO, an automatic prompt engineering approach tailored for generative tasks. It discovers aspects to critique generated text and write suggestions for more effective prompt revision.

2) We conduct comprehensive experiments across multiple LLMs and datasets, demonstrating the effectiveness and robustness of our method. We show an overall 3-4% improvement on ROUGE scores with qualitative verification from human evaluation. CriSPO also obtained consistent improvements on various QA tasks.

3) We propose AST that enables prompt tuning for multiple metrics. We show that CriSPO with AST can jointly optimize AlignScore (Zha et al. 2023) for faithfulness and ROUGE for reference similarity.

## 2 Related Work

There is an increasing effort in the literature to explore gradient-free automatic prompt engineering methods with off-the-shelf LLMs. The focus of these approaches is to find a good search algorithm for better prompt candidates to solve discriminative tasks. Earlier studies have employed conventional paraphrasing methods for prompt generation through editing phrases (Prasad et al. 2023) or back translation (Xu

et al. 2022). More recently, LLMs themselves have been used to sample prompt candidates. Zhou et al. (2022) proposed Automatic Prompt Engineering (APE) which iteratively prompts an LLM to generate semantically similar variations of the locally best prompt. Pryzant et al. (2023) add verbal feedback based on error examples to propose better prompts in terms of accuracy. Concurrently, Sordoni et al. (2024) learn prompts with variational inference by considering their outputs as latent variables. Later on, Yang et al. (2023) propose OPRO to improve over them by incorporating the history of past prompts with their scores which stabilizes optimization. More structured prompts have also been explored by imposing expert-level planning (Wang et al. 2023). In a parallel thread, Fernando et al. (2023) and Guo et al. (2023) were inspired by evolutionary algorithms to perform mutation operations for prompt generation. All of the existing approaches have mostly been designed to target classification tasks using a single metric. Comparing to the existing studies, our proposed method specifically targets the unique challenges in text generation and approaches the prompt optimization problem in a multi-aspect and multi-metric fashion. For practitioners, Khattab et al. (2023) design DSPy framework to build and optimize complex LLM pipelines in a programmatic fashion. TextGrad (Yuksekgonul et al. 2024) further generalizes optimization to text beyond prompt. Our CriSPO can be used as a powerful optimizer in these frameworks.

Our approach is also inspired by recent studies on using LLMs to automatically correct its output (Pan et al. 2023; Madaan et al. 2024). Gero et al. (2023) apply multiple self-reflection steps to improve the performance of information extraction. Yan et al. (2024) use CoT to generate structured comparison and preferences for two model outputs. Shinn et al. (2023) argue the importance of the self-reflection history and propose reflexion agent to provide verbal feedback on past trials for better decision in the next trials. It is important to notice that these self-reflection studies are strictly speaking not automatic prompt engineering approaches as these studies optimize output revision rather than directly on the prompts. CriSPO, however, automatically reflects on the design of the prompt and uses these past reflections to revise the prompts.

## 3 Method

**Problem Formulation:** In a text generation task, let  $\mathcal{D}_{\text{trn}} = \{(x_i, y_i)\}_{i=1\dots n}$  be the training set, with a development set  $\mathcal{D}_{\text{dev}}$  and a test set  $\mathcal{D}_{\text{tst}}$ . Here,  $x$  represents the input data, and  $y$  is the corresponding ground truth reference. A **task prompt**  $p$  comprises instructions that, when filled with input  $x$ , are fed to a black-box API  $\text{LLM}_{\text{task}}$ <sup>1</sup> to generate a completion  $\hat{y} = \text{LLM}_{\text{task}}(p, x)$ . The goal is to optimize  $p$  using  $\mathcal{D}_{\text{trn}}$  and  $\mathcal{D}_{\text{dev}}$  to identify an optimal prompt  $p^*$  that maximizes performance on one or more evaluation metrics on  $\mathcal{D}_{\text{tst}}$ .

**CriSPO Overview:** CriSPO is an automatic prompt optimization algorithm designed to iteratively refine a task prompt  $p$  from an initial seed prompt  $p_0$  to the optimum:  $p^* \leftarrow \mathcal{F}(p_0)$ . In each iteration  $t$ , we conduct the following steps:

<sup>1</sup>We use notations  $\text{LLM}_{\text{task}}$ ,  $\text{LLM}_{\text{crit}}$ ,  $\text{LLM}_{\text{opti}}$  for clarity. Though they share the same underlying LLM unless specified otherwise.

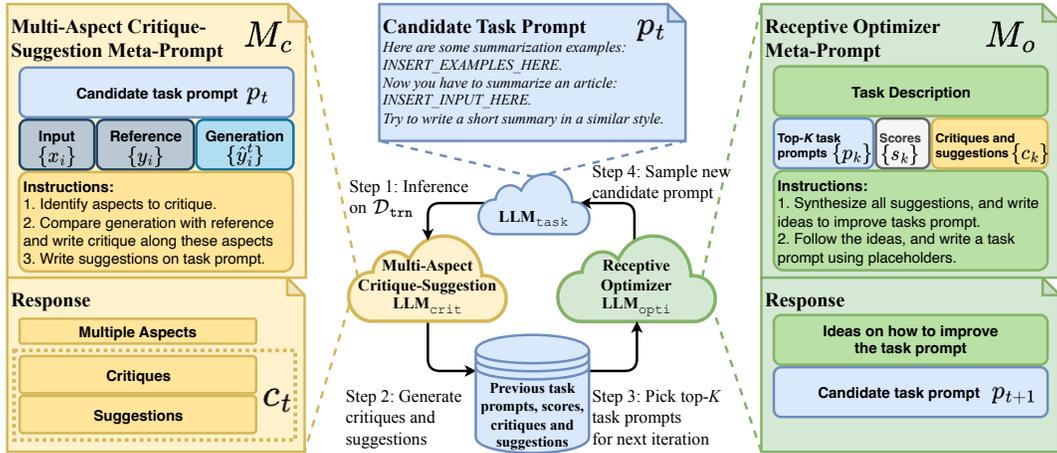


Figure 1: The CriSPO workflow for text generation tasks. In each iteration, a candidate task prompt  $p_t$  is applied to  $\mathcal{D}_{\text{trn}}$  (step 1) and evaluated using a multi-aspect critique-suggestion meta-prompt  $M_c$  (step 2). We select top- $K$  previously sampled task prompts (step 3) and use a receptive optimizer  $M_o$  to generate the next candidate  $p_{t+1}$  (step 4). The automatic optimization loop runs multiple iterations, while the best task prompt is selected based on performance on  $\mathcal{D}_{\text{dev}}$ .

- **Evaluate on  $\mathcal{D}_{\text{trn}}$ :** Apply the candidate prompt  $p_t$  on  $\mathcal{D}_{\text{trn}}$ , call  $\text{LLM}_{\text{task}}$  to generate outputs  $\{\hat{y}_i^t\}_{i=1\dots n}$  and compute a primary metric  $s_t$ , which can be a single metric or an aggregation of multiple metrics.
- **Generate Critiques and Suggestions:** Apply the **multi-aspect critique-suggestion** meta-prompt  $M_c$  and call  $\text{LLM}_{\text{crit}}$  to compare  $\{\hat{y}_i^t\}_{i=1\dots n}$  and  $\{y_i^t\}_{i=1\dots n}$  and generate critiques and suggestions  $c_t$  (see Section 3.1).
- **Generate a Candidate Task Prompt:** Select the top- $K$  task prompts from previous iterations based on the primary metric, and insert the corresponding  $K$  triples  $\{(p_k, s_k, c_k)\}$  into the **receptive optimizer** meta-prompt  $M_o$ . Then call  $\text{LLM}_{\text{opti}}$  to generate the next candidate prompt  $p_{t+1}$  (see Section 3.2).

We evaluate the current prompt  $p_t$  on  $\mathcal{D}_{\text{dev}}$  and select  $p^*$  based on the primary metric. Upon reaching the maximum number of iterations, we apply an optional AST to enhance performance on secondary metrics on  $p^*$  (see Section 3.3). Figure 1 demonstrates the workflow of CriSPO on summarization tasks. Table 8 (in Appendix) shows a concrete working example of CriSPO.

### 3.1 Multi-Aspect Critiques and Suggestions

Given a prompt  $p_t$  and its outputs  $\{\hat{y}_i^t\}$  on  $\mathcal{D}_{\text{trn}}$ , we design a multi-aspect critique-suggestion meta-prompt  $M_c$  to identify *critiques* – flaws of the generated outputs across multiple aspects, and *suggestions* – specific edits on the task prompt to rectify each flaw.

**Constructive critiques with spontaneous dimension discovery:** In  $M_c$ , we first instruct  $\text{LLM}_{\text{crit}}$  to generate several task-specific and iteration-specific aspects for a given batch of outputs from the current  $p_t$ . This approach ensures that as task prompts evolve across iterations, the focus remains on relevant aspects, addressing specific issues that

arise. Figure 2 illustrates the aspects discovered during optimization. For each aspect,  $M_c$  instructs  $\text{LLM}_{\text{crit}}$  to generate a critique highlighting potential problems of the outputs generated with  $p_t$  on the batch.

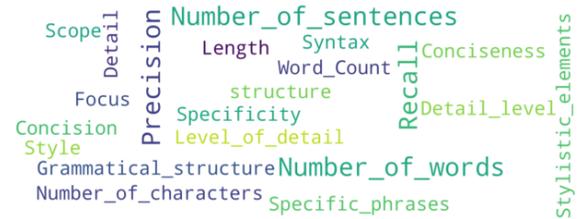


Figure 2: A word cloud showing the different aspects identified by CriSPO when comparing generations and references.

**Multi-aspect suggestions:** In line with each critique, a corresponding suggestion is made by  $\text{LLM}_{\text{crit}}$  to edit  $p_t$ . As opposed to Pryzant et al. (2023), we decoupled the edit suggestion module from the new prompt generation process. Rather than generating a new prompt with each suggestion, we pack a history of critiques and suggestions into the receptive optimizer for generating the next prompt, enabling more stable optimization over the infinite search space.

Our  $M_c$  is implemented in a single CoT meta-prompt which generates, dimensions, critiques and suggestions in one single LLM call, specifically

$$c_t = \text{LLM}_{\text{crit}}(M_c, p_t, (x_i, y_i, \hat{y}_i^t)_{i=1\dots n}).$$

$M_c$  for different LLMs and tasks are shown in Appendix H.

### 3.2 Receptive Prompt Optimizer

Our receptive prompt optimizer meta-prompt  $M_o$  improves over the OPRO optimizer meta-prompt (Yang et al. 2023) by enriching its optimization trajectory  $\{(p_k, s_k)\}$  with past

Dataset	LLM	Manual						Automatic Prompt Engineering								
		0-shot			3-shot*			OPRO			CriSPO			CriSPO 3-shot*		
		R1	R2	RL	R1	R2	RL	R1	R2	RL	R1	R2	RL	R1	R2	RL
CNN	Claude In.	37.5	12.5	22.6	40.4	14.8	24.8	39.5	14.3	24.5	40.1	15.7	26.1	<b>42.1</b>	<b>17.0</b>	<b>27.4</b>
	Claude3	38.8	14.4	24.0	40.3	15.4	25.2	39.7	15.1	5.1	<b>42.2</b>	<b>17.3</b>	<b>27.9</b>	41.6	16.3	27.1
	Mistral 7B	30.9	11.0	20.4	30.7	10.6	20.1	36.5	<b>14.4</b>	23.0	38.5	14.3	23.9	<b>38.5</b>	14.3	<b>24.1</b>
	Llama3 8B	37.9	14.4	23.8				39.1	15.2	24.6 <sup>#</sup>	<b>41.5</b>	<b>16.3</b>	<b>26.5<sup>#</sup></b>			
MeetingBank	Claude In.	30.7	11.6	20.5	34.2	17.3	25.5	39.0	20.3	29.7	41.4	23.7	33.1	<b>50.1</b>	<b>35.4</b>	<b>44.4</b>
	Claude3	31.2	14.2	22.3	37.5	22.0	29.5	41.5	21.8	32.0	47.4	32.5	40.9	<b>58.5</b>	<b>46.5</b>	<b>54.1</b>
	Mistral 7B	26.0	11.5	18.5	31.3	14.8	22.7	33.9	15.4	24.2	<b>39.1</b>	<b>19.5</b>	<b>29.3</b>	35.2	16.7	26.1
	Llama3 8B	31.4	14.6	22.6				40.2	22.3	31.5 <sup>#</sup>	<b>44.7</b>	<b>27.6</b>	<b>36.8<sup>#</sup></b>			
SAMSum	Claude In.	33.9	11.7	25.6	37.8	14.3	28.8	38.1	13.4	28.7	44.4	16.9	34.3	<b>45.7</b>	<b>18.7</b>	<b>36.2</b>
	Claude3	35.8	12.7	27.0	41.1	16.6	31.3	39.0	14.7	30.1	43.4	17.1	34.3	<b>47.2</b>	<b>20.8</b>	<b>38.2</b>
	Mistral 7B	32.0	10.2	24.1	39.5	14.1	30.3	37.9	13.6	29.0	37.6	12.4	28.4	<b>40.0</b>	<b>14.2</b>	<b>30.8</b>
	Llama3 8B	35.7	12.3	27.1				39.3	14.7	30.0 <sup>#</sup>	<b>44.8</b>	<b>18.8</b>	<b>35.4<sup>#</sup></b>			
ACI-Bench	Claude In.	43.8	16.9	26.1	51.5	23.6	33.5	45.2	16.3	25.5	53.0	19.7	26.8	<b>58.2</b>	<b>26.7</b>	<b>35.3</b>
	Claude3	47.3	20.3	29.3	59.1	30.1	38.6	48.8	20.1	29.5	54.0	21.4	30.3	<b>63.1</b>	<b>32.5</b>	<b>41.0</b>
	Mistral 7B	47.8	17.7	25.4	48.4	<b>19.2</b>	<b>28.1</b>	45.1	17.0	25.2	50.2	18.2	25.6	<b>50.3</b>	18.7	26.2
	Llama3 8B	50.5	19.8	27.7				54.2	22.0	29.3 <sup>#</sup>	<b>56.2</b>	<b>22.8</b>	<b>29.9<sup>#</sup></b>			
Average	Claude In.	36.5	13.2	23.7	41.0	17.5	28.2	40.4	16.1	27.1	44.7	19.0	30.1	<b>49.0</b>	<b>24.4</b>	<b>35.8</b>
	Claude3	38.3	15.4	25.6	44.5	21.0	31.2	42.2	17.9	29.2	46.8	22.1	33.3	<b>52.6</b>	<b>29.0</b>	<b>40.1</b>
	Mistral 7B	34.2	12.6	22.1	37.5	14.7	25.3	38.4	15.1	25.4	<b>41.4</b>	<b>16.1</b>	<b>26.8</b>	41.0	16.0	<b>26.8</b>
	Llama3 8B	38.9	15.3	25.3				43.2	18.6	28.8	<b>46.8</b>	<b>21.4</b>	<b>32.2</b>			

Table 1: Comparing CriSPO with manual prompts and OPRO on representative summarization benchmarks. Averaged R1/R2/RL (i.e. ROUGE-1/2/L) are reported across 3 runs. 3-shot\*: 3-shot ICL with example selection. Claude3: Claude3 Sonnet. Claude In.: Claude Instant. Llama3 results marked with (#) are using Claude3 Sonnet as the optimizer. 3-shot\* results are empty for Llama3 due to its limited context window. Standard deviation and SOTA results are in Appendix M. Claude3 Sonnet achieves new SOTA ROUGE-1 performance on ACI-Bench.

critiques and suggestions  $c_k$ . Thus, ours samples candidate prompts for the next iteration conditioned on an enriched optimization trajectory:

$$p_{t+1} = \text{LLM}_{\text{opti}}(M_o, \{(p_k, s_k, c_k)\}).$$

Specifically, we enhance the OPRO optimizer module with the following three improvements to better utilize critiques and suggestions for achieving stronger guidance and better exploration. See Appendix I for all  $M_o$  by LLMs and tasks.

**Enriched optimization trajectory:** The critiques and suggestions generated in Section 3.1 are used in an enriched optimization trajectory to propose new prompts via an OPRO-style optimizer. Specifically, our enriched optimization trajectory includes the top- $K$  best-performing past prompts  $\{p_k\}$ , their scores  $\{s_k\}$ , critiques and suggestions  $\{c_k\}$ , sorted in the ascending order by scores. Including critiques and suggestions in the optimization trajectory allows the LLM to avoid common limitations and identify common strengths from the past prompts for stable optimization.

**Chain-of-thought:** After enriching the optimization trajectory, we also apply CoT to the optimization process. Specifically,  $\text{LLM}_{\text{opti}}$  is explicitly asked to first compare high-score prompts to low-score ones, and then elicit general ideas and learnings, and finally draft a new and better prompt. CoT further ensures the optimizer to harness collective strength from

the history and identify a promising path through comparing the divergent past prompts.

**Flexible task prompt template:** Instead of only tuning instruction text and fixing the input position as in existing approaches such as OPRO, CriSPO optimizes the task prompt structure using a template that can freely and naturally move around input and instruction in the prompt. It uses placeholders for the input and any external data. For example, we instruct LLM to generate example placeholder `INSERT_EXAMPLES_HERE` to indicate the position of ICL examples. In Retrieval-Augmented Generation (RAG) settings, we introduce a context placeholder `INSERT_CONTEXT_HERE` which will be replaced by the retrieved context for each question. When filled the placeholders with proper data, the task prompt clearly organized all the information to help  $\text{LLM}_{\text{task}}$  better solve the task.

### 3.3 Multi-Metric Automatic Suffix Tuning

Using components in Section 3.1 and 3.2, CriSPO is ready to optimize a primary metric. To benefit from more teaching signals, e.g., completeness and faithfulness, here we extend CriSPO to multi-metric optimization by proposing a novel multi-metric learning extension named as AST.

In AST, we propose to optimize a suffix postscript  $\sigma$  appended to  $p^*$ , which has already been trained on certain metrics.  $p^*$  will remain fixed throughout the whole tuning

process for a new metric to preserve most of its performance on existing metrics.  $p^*$  is extended with an additional suffix  $\sigma^* \leftarrow \mathcal{F}(\sigma_0)$ , which serves as a postscript to steer the LLM toward the new metric and remedy any potential regression in performance on existing metrics. Specifically, we provide both the main prompt  $p^*$  and each suffix  $\sigma_t$  in the meta-prompts while asking the LLM to critique or refine only the suffix. To ensure we maintain existing metrics while improving on the additional metric, we take inspirations from the balance terms of loss functions in multi-task learning (He and Choi 2021) and compute an aggregated score across the multiple metrics. Since the score of each metric is on different scales and hard to estimate before training, we propose to use the average ranking of each metric as the ultimate basis to score prompt candidates in the meta-prompt.

## 4 Main Experiments

### 4.1 Experiment Setup

**Datasets** We select a diverse range of 4 summarization tasks including conventional document summarization tasks such as CNN daily mail (Hermann et al. 2015) (news headline summarization), and also conversation summarization tasks such as SAMSum (Gliwa et al. 2019), MeetingBank (Hu et al. 2023). In addition, we test on a medical-domain clinical note summarization task, ACI-Bench (Yim et al. 2023). Detailed data setup can be found in Appendix C. These tasks cover various lengths, domains and styles as summarized in Table 9. We report ROUGE-1/2/L F-measure (Lin 2004)<sup>2</sup> to measure output similarity to the references.

**LLMs and Baselines** We test our approach on state-of-the-art LLMs including proprietary models: Claude Instant (Anthropic 2023), Claude3 Sonnet (Anthropic 2024), and open-source LLMs: Mistral 7B (Jiang et al. 2023) and Llama3 8B (MetaAI 2024). We use the same LLM for all the 3 CriSPO modules: task inference, critique-suggestion and receptive optimization, apart from the Llama3 setup. Specific hyperparameters with ablations are detailed in Appendix D.

Our baseline methods include manual prompts with zero/few-shot ICL. These manual prompts are carefully tuned for each task to incorporate length constraints and task guidelines, and therefore establish a high bar of performance from manual prompt engineering (Appendix J). Given there are no existing automatic prompting results for text generation, we adapted OPRO (Yang et al. 2023), a competitive established approach, on our selected tasks. We use the same hyperparameter setup in OPRO and CriSPO for fair comparison.

### 4.2 Main Results

As shown in Table 1, across all the tasks and LLMs, CriSPO consistently improves over 0-shot manual prompt and OPRO baselines. Overall, there are approximately 3-4 point improvements for all LLMs. Even the strong state-of-the-art Claude3 Sonnet model can still greatly benefit from CriSPO. The consistent improvement shows CriSPO is a more effective search

method than existing method (OPRO) to unlock the full potential of these LLMs, and offers an alternative solution to the more labour-intensive manual prompt engineering.

Additionally, we found examples to be helpful as adding 3-shot ICL significantly improves the performance. Owing to the versatile template in CriSPO, we can easily integrate examples and we show CriSPO 3-shot can further boost performance over CriSPO and achieves the best performance in most setups. It is also worth noticing that the vanilla CriSPO w/o ICL can match or even outperform manual prompt with 3-shot in most datasets and setups, reducing latency and cost.

### 4.3 Ablating Key Ingredients

Table 2 shows the ablation results of CriSPO with Claude Instant on SAMSum dataset. We observed that the three key components in our approach, including flexible template, critique and step-by-step CoT optimization, are essential for achieving optimal performance. Removing any of these components leads to a decrease in performance. Removing critique-suggestion module and CoT optimization altogether leads to a 5 point decrease, similar to OPRO performance. This indicates these two elements are essential to the success of CriSPO and flexible template is only effective when being added on top of these two elements.

Method	Crit-Sugg	CoT	Template	avg	std
CriSPO	✓	✓	✓	44.4	1.9
	✗	✓	✓	42.8	0.8
	✓	✗	✓	43.9	0.3
	✓	✓	✗	42.2	1.6
	✗	✗	✓	37.4	3.4
OPRO	✗	✗	✗	38.1	1.3
Method	Changing multi-aspect			avg	std
CriSPO	free multi-aspects			44.4	1.9
	no multi-aspects			41.1	0.9
	pre-defined multi-aspects			44.5	0.7

Table 2: Ablation studies of CriSPO on the SAMSum dataset with Claude Instant. We report average and std deviation of the ROUGE-1 F results across 3 runs.

The key novelty in our proposed novel critique-suggestion strategy in CriSPO is that it has multi-aspect: i.e. the LLM will generate multi-aspect comparison without enforcing pre-defined aspects. To understand the effect of the multi-aspect critique-suggestion, we provide two alternative baselines: 1. no multi-aspect: we ask LLM to compare predictions and references in general with no explicit requirement for generating critique and suggestions along multiple dimensions/aspects. This is in line with the approach adopted by Pryzant et al. (2023). 2. predefined aspects: we carefully design dimensions potentially helpful for the summarization task and include verbosity, comprehensiveness, precision and style along with their definitions (Appendix L). The no multi-aspect critique-suggestion baseline performs significantly worse, lacking critical and targeted suggestions due to its tendency to be too general. The predefined multi-aspect approach is as effective as CriSPO but we see no significant improvement from explicit definitions of dimensions. This is because the critique

<sup>2</sup>We report additional metrics including AlignScore (Zha et al. 2023) and BertScore (Zhang et al. 2019) in Appendix M.

LLM in CriSPO is already able to identify relevant dimensions (such as completeness, verbosity etc. as in Table 8) for each iteration without explicit guidance.

#### 4.4 Qualitative Analysis and Human Evaluation

To qualitatively compare CriSPO outputs with the baselines, we conducted human evaluation on 20 examples from the SAMSum testset. We follow the procedure from Liu et al. (2023) where the reference summaries are split into atomic content units and annotators mark them as either present or missing in the prediction summary. In total, we collected 300 annotations (100 annotations  $\times$  3 annotators). A final normalized recall score is computed with a length penalty, which indicates how similar the prediction summary is to the reference summary. In our experiment, we asked three annotators with postgraduate degrees to independently annotate the summaries with blinded setup. The inter-annotator agreement is “almost perfect” (0.8679 Fleiss kappa). We then took the majority vote, and calculated the final normalized recall score (human rating) using a de-correlated length penalty.

As shown in Table 3, CriSPO achieves the highest rating according to our human evaluation. Table 4 shows qualitative examples where prompts found by CriSPO better capture the style of the reference summaries in terms of length, what to focus on and what to skip. CriSPO outputs also look the most similar to the references, especially in terms of being as concise as the reference while covering all the key details.

	Manual	OPRO	CriSPO
<b>Human Rating</b>	0.58	0.59	<b>0.63</b>

Table 3: Human evaluation on sampled SAMSum test.

**OPRO [Best Prompt]:** Generate a **one to two sentence** summary within the  $\langle$ summary $\rangle$  tags that concisely describes the key details of the conversation and **any conclusions** reached. INPUT\_DOC

**CriSPO [Best Prompt]:** The text below contains a discussion expressing several **key facts and events**. Your concise **1-sentence** summary should relate only the **2 most** important pieces of information stated, **without assumptions or extra context**. INPUT\_DOC Write the summary within  $\langle$ summary $\rangle$  tags.

**OPRO [Example Output]:** Ralph asked Andrew if he heard a Polish joke, then told a joke about sinking a Polish battleship by putting it in water. Andrew responded that the joke was terrible and so unfunny that it made his mouth dry, requiring a sip of water.

**CriSPO [Example Output]:** Ralph tells Andrew a Polish battleship joke that Andrew finds unfunny.

**[Reference]:** Ralph told Andrew a joke.

Table 4: Qualitatively comparing prompts found by OPRO and CriSPO on SAMSum. CriSPO is able to find a prompt to generate output more similar to the reference’s concise style.

#### 4.5 Quantitative Analysis of Prompt Diversity

To verify that our design in Section 3 leads to a better exploration of the solution space, we quantitatively analyze the diversity of prompts found by CriSPO and OPRO (same hyperparameters, Section 4.1) on the summarization datasets. We measure 4 aggregated properties on all task prompts explored by each method during optimization: length (number of words), vocabulary size (number of unique words used), and pairwise ROUGE-L/semantic similarity. For pairwise semantic similarity, we employ Sentence Transformers (Reimers and Gurevych 2019) to obtain their embeddings and cosine distances.

As shown in Table 5, CriSPO prompts demonstrate larger variations in length and vocabulary while being less similar in lexicons and semantics, indicating its strength in exploring a larger space. We also provide a visualization of the prompts found by OPRO and CriSPO in Appendix F.

Dataset	Length $\uparrow$	Vocab $\uparrow$	ROUGE-L $\downarrow$	Cosine $\downarrow$
<b>CNN</b>				
OPRO	41 $\pm$ 6	36 $\pm$ 5	57.5	0.93
CriSPO	<b>149<math>\pm</math>24</b>	<b>96<math>\pm</math>12</b>	<b>50.3</b>	<b>0.90</b>
<b>MeetingBank</b>				
OPRO	31 $\pm$ 5	28 $\pm$ 4	44.9	0.84
CriSPO	<b>216<math>\pm</math>41</b>	<b>135<math>\pm</math>19</b>	<b>39.7</b>	<b>0.80</b>
<b>SAMSum</b>				
OPRO	34 $\pm$ 6	30 $\pm$ 5	57.0	0.94
CriSPO	<b>172<math>\pm</math>22</b>	<b>112<math>\pm</math>12</b>	<b>46.0</b>	<b>0.88</b>
<b>ACI-Bench</b>				
OPRO	58 $\pm$ 11	46 $\pm$ 8	62.7	0.95
CriSPO	<b>247<math>\pm</math>40</b>	<b>117<math>\pm</math>13</b>	<b>54.3</b>	<b>0.93</b>

Table 5: Prompt diversity on 4 summarization datasets.

### 5 Extension with Multi-Metric Optimization

**AST Setup** In this experiment, we extend CriSPO with our proposed AST to optimize multiple metrics simultaneously. Specifically, we take the best prompts optimized for ROUGE-1 F-measure from CriSPO with Claude Instant as the seed main prompt  $p^*$ . We employ AST to optimize AlignScore (Zha et al. 2023) starting from a simple seed suffix  $\sigma_0$ : “Every word of your summary must be faithful to the input/conversation” across all datasets. The AlignScore between the input text and the output summary is used as a signal reflecting the faithfulness. With regard to baselines, we report the initial performance in ROUGE-1 F-measure and AlignScore of the seed main prompt w/ and w/o the seed suffix. We also provide a strong baseline to tune both the main prompt and its suffix together (full tuning) rather than only the suffix in AST.

**Results** The results for multi-metric optimization are presented in Table 6. On all datasets, our AST is able to optimize the new metric AlignScore with a negligible or zero regression on the existing metric ROUGE, meaning that AST can reduce LLM hallucination while maintaining relevancy in the output. In particular, AST dramatically improves AlignScore by 11.7 points on CNN. Across tasks, AST is the most

effective approach to improve AlignScore while maintaining ROUGE. Among all methods, AST is the only one that brings consistent improvement on AlignScore for every task, and achieves the best average overall improvement (by 4.3). The main prompt w/ suffix seed prompt slightly improves AlignScore (by 1.2) and the full-tuning baseline only meaningfully improves AlignScore on CNN and the overall improvement is marginal (by 0.7). The superiority of AST shows that it can robustly optimize multiple metrics across various domains.

Dataset	Seed		CriSPO: $\mathcal{F}(\cdot)$	
	main $p^*$	w/ suffix $p^* + \sigma_0$	full $\mathcal{F}(p^* + \sigma_0)$	w/ AST $p^* + \mathcal{F}(\sigma_0)$
<b>CNN</b>				
ROUGE-1	<b>40.7</b>	40.6	40.6	40.4
AlignScore	66.5	69.5( $\uparrow$ 3.0)	69.6( $\uparrow$ 3.1)	<b>78.1(<math>\uparrow</math>11.7)</b>
<b>MeetingBank</b>				
ROUGE-1	39.6	<b>39.9</b>	39.4	39.7
AlignScore	43.6	43.7	43.8	<b>44.4(<math>\uparrow</math>0.9)</b>
<b>SAMSum</b>				
ROUGE-1	45.5	<b>45.9</b>	45.8	45.1
AlignScore	87.2	86.6( $\downarrow$ 0.6)	86.6( $\downarrow$ 0.6)	<b>88.6(<math>\uparrow</math>1.4)</b>
<b>ACI-Bench</b>				
ROUGE-1	54.4	<b>55.2(<math>\uparrow</math>0.8)</b>	54.5	54.3
AlignScore	66.7	69.0( $\uparrow$ 2.3)	66.5	<b>70.0(<math>\uparrow</math>3.4)</b>
<b>Average</b>				
ROUGE-1	45.1	45.4	45.1	44.9
AlignScore	66.0	67.2( $\uparrow$ 1.2)	66.7( $\uparrow$ 0.7)	<b>70.3(<math>\uparrow</math>4.3)</b>

Table 6: Claude Instant multi-metric results on summarization tasks. In seed block, main and w/ suffix refers to the ROUGE-1-optimized prompt  $p^*$  and its concatenation with the manual suffix  $\sigma_0$  respectively. In CriSPO block, full and w/ AST refers to the full-prompt tuning baseline and AST in Section 3.3 respectively. We show in the parentheses the absolute differences with main  $p^*$  only when  $> 0.5$ . Here we use CriSPO  $\mathcal{F}(\cdot)$  to optimize an aggregation of two metrics.

## 6 Generalization to Other Tasks

To confirm its generalizability, in this section, we apply CriSPO to extractive, abstractive and multi-choice QA tasks.

**Datasets** We benchmark CriSPO on 5 commonly used QA datasets, including 1) Wikipedia-based QA: Natural Questions (Kwiatkowski et al. 2019), TriviaQA (Joshi et al. 2017), Squad (Rajpurkar et al. 2016) 2) story-based abstractive reading comprehension: NarrativeQA (Kočíský et al. 2018) and 3) medical domain multiple-choice QA: MedMCQA (Pal, Umapathi, and Sankarasubbu 2022). For Natural Questions and TrivialQA, we also incorporate the RAG setup to optimize the prompt template with inserted pre-retrieved contexts from each dataset. We retrieved the Wikipedia pages following Izacard and Grave (2021). For NarrativeQA, we use summaries as contexts. For MedMCQA, we cast it to text generation by eliciting reasoning before the final answer. Following the conventions, we report Exact Match for Natural Questions and TriviaQA, F1 for Squad, ROUGE-L for NarrativeQA, accuracy for MedMCQA. For efficiency, we only

used a small fraction of the train and dev set for the experiments. The specific data settings are listed in Appendix C.

**Results** Similar to summarization tasks, we observe CriSPO significantly outperforms the manual prompt and OPRO baseline in various QA datasets as shown in Table 7. For NarrativeQA, CriSPO brings massive improvement (+10 ROUGE-L) compared with baselines, achieving the new SOTA performance. For Natural Questions and TrivialQA, CriSPO has no issue incorporating the RAG setup and achieving consistent improvement over the manual prompt and OPRO. Surprisingly, CriSPO even outperforms OPRO on MedMCQA despite it is not designed for classification tasks.

Task	Claude	Manual		Automatic Prompt Engineering		
		0-shot	64*	OPRO	CriSPO	CriSPO 64*
NQ	Instant	34.0	33.4	8.0	36.5	<b>37.8</b>
	Sonnet	26.6	32.0	6.7	38.3	<b>38.7</b>
T-QA	Instant	58.6	59.2	53.7	66.3	<b>67.5</b>
	Sonnet	58.4	65.0	41.8	70.6	<b>72.1</b>
		0-shot	5*	OPRO	CriSPO	CriSPO 5*
Squad	Instant	79.5	82.5	78.5	87.8	<b>89.4</b>
	Sonnet	76.1	83.2	76.4	85.3	<b>87.9</b>
NarQA	Instant	64.2	67.0	59.4	75.1	<b>76.1</b>
	Sonnet	64.0	66.7	58.6	<b>76.2</b>	75.2
Med-MCQA	Instant	49.2	53.8	50.5	52.3	<b>54.4</b>
	Sonnet	49.8	54.4	57.7	<b>57.9</b>	57.4

Table 7: Comparing CriSPO with manual prompts and competitive automatic prompt engineering baseline OPRO on representative QA benchmarks. We report Exact matching for NQ (Natural Questions) and TQA (TrivialQA). We report F1 for Squad, Rouge-L for NarQA (NarrativeQA), and accuracy for MedMCQA.  $k^*$ :  $k$ -shot ICL with example selection. Standard deviations can be found in Table 17.

## 7 Conclusion

In this paper, we tackle the challenging problem of automatic prompt engineering for text generation. We propose CriSPO, a multi-aspect critique-suggestion guided optimizer augmented with enriched trajectory, CoT and flexible template. Our experiments show multi-aspect critique-suggestion is critical for finding good task prompts. Overall, CriSPO achieves 3-4% ROUGE score improvement and 4-5% human rating increase compared to baseline methods for summarization, and significant improvement for QA. We also show that CriSPO can effectively optimize multiple metrics through a novel suffix tuning extension AST, and incorporate ICL and RAG with flexible prompt templates. Ablation studies confirm the effectiveness of all CriSPO components. Human evaluation and quantitative analysis show CriSPO encourages more effective prompt exploration and the optimized prompts can better capture task requirements.

## Limitations

See Appendix B for a detailed limitations discussion.

## References

- Anthropic. 2023. Claude instant model 1.2. Accessed: 2024-06-13.
- Anthropic, A. 2024. The Claude 3 model family: Opus, sonnet, haiku. *Claude-3 Model Card*.
- Augenstein, I.; Baldwin, T.; Cha, M.; Chakraborty, T.; Ciampaglia, G. L.; Corney, D.; DiResta, R.; Ferrara, E.; Hale, S.; Halevy, A.; et al. 2023. Factuality challenges in the era of large language models. *arXiv preprint arXiv:2310.05189*.
- Brown, T.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J. D.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33: 1877–1901.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In Burstein, J.; Doran, C.; and Solorio, T., eds., *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, 4171–4186. Minneapolis, Minnesota: Association for Computational Linguistics.
- Elangovan, A.; Liu, L.; Xu, L.; Bodapati, S.; and Roth, D. 2024. ConSiDERS-The-Human Evaluation Framework: Rethinking Human Evaluation for Generative Large Language Models. *arXiv preprint arXiv:2405.18638*.
- Fabbri, A. R.; Kryściński, W.; McCann, B.; Xiong, C.; Socher, R.; and Radev, D. 2021. SummEval: Re-evaluating Summarization Evaluation. *Transactions of the Association for Computational Linguistics*, 9: 391–409.
- Fernando, C.; Banarse, D.; Michalewski, H.; Osindero, S.; and Rocktäschel, T. 2023. Promptbreeder: Self-referential self-improvement via prompt evolution. *arXiv preprint arXiv:2309.16797*.
- Gao, M.; and Wan, X. 2022. DialSummEval: Revisiting Summarization Evaluation for Dialogues. In Carpuat, M.; de Marneffe, M.-C.; and Meza Ruiz, I. V., eds., *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 5693–5709. Seattle, United States: Association for Computational Linguistics.
- Gardent, C.; Shimorina, A.; Narayan, S.; and Perez-Beltrachini, L. 2017. Creating training corpora for nlg micro-planning. In *55th Annual Meeting of the Association for Computational Linguistics, ACL 2017*, 179–188. Association for Computational Linguistics (ACL).
- Gero, Z.; Singh, C.; Cheng, H.; Naumann, T.; Galley, M.; Gao, J.; and Poon, H. 2023. Self-verification improves few-shot clinical information extraction. In *ICML 3rd Workshop on Interpretable Machine Learning in Healthcare (IMLH)*.
- Gliwa, B.; Mochol, I.; Biesek, M.; and Wawer, A. 2019. SAMSum Corpus: A Human-annotated Dialogue Dataset for Abstractive Summarization. In Wang, L.; Cheung, J. C. K.; Carenini, G.; and Liu, F., eds., *Proceedings of the 2nd Workshop on New Frontiers in Summarization*, 70–79. Hong Kong, China: Association for Computational Linguistics.
- Guo, Q.; Wang, R.; Guo, J.; Li, B.; Song, K.; Tan, X.; Liu, G.; Bian, J.; and Yang, Y. 2023. Connecting Large Language Models with Evolutionary Algorithms Yields Powerful Prompt Optimizers. In *The Twelfth International Conference on Learning Representations*.
- He, H.; and Choi, J. D. 2021. The Stem Cell Hypothesis: Dilemma behind Multi-Task Learning with Transformer Encoders. In Moens, M.-F.; Huang, X.; Specia, L.; and Yih, S. W.-t., eds., *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, 5555–5577. Online and Punta Cana, Dominican Republic: Association for Computational Linguistics.
- Hermann, K. M.; Kocisky, T.; Grefenstette, E.; Espeholt, L.; Kay, W.; Suleyman, M.; and Blunsom, P. 2015. Teaching machines to read and comprehend. *Advances in neural information processing systems*, 28.
- Hu, Y.; Ganter, T.; Deilamsalehy, H.; Dernoncourt, F.; Foroosh, H.; and Liu, F. 2023. MeetingBank: A Benchmark Dataset for Meeting Summarization. In Rogers, A.; Boyd-Graber, J.; and Okazaki, N., eds., *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 16409–16423. Toronto, Canada: Association for Computational Linguistics.
- Izacard, G.; and Grave, E. 2021. Leveraging Passage Retrieval with Generative Models for Open Domain Question Answering. In Merlo, P.; Tiedemann, J.; and Tsarfaty, R., eds., *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, 874–880. Online: Association for Computational Linguistics.
- Izacard, G.; Lewis, P.; Lomeli, M.; Hosseini, L.; Petroni, F.; Schick, T.; Dwivedi-Yu, J.; Joulin, A.; Riedel, S.; and Grave, E. 2023. Atlas: Few-shot learning with retrieval augmented language models. *Journal of Machine Learning Research*, 24(251): 1–43.
- Jiang, A. Q.; Sablayrolles, A.; Mensch, A.; Bamford, C.; Singh Chaplot, D.; de las Casas, D.; Bressand, F.; Lengyel, G.; Lample, G.; Saulnier, L.; et al. 2023. Mistral 7B. *arXiv e-prints*, arXiv–2310.
- Joshi, M.; Choi, E.; Weld, D.; and Zettlemoyer, L. 2017. TriviaQA: A Large Scale Distantly Supervised Challenge Dataset for Reading Comprehension. In Barzilay, R.; and Kan, M.-Y., eds., *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 1601–1611. Vancouver, Canada: Association for Computational Linguistics.
- Khattab, O.; Singhvi, A.; Maheshwari, P.; Zhang, Z.; Santhanam, K.; Vardhamanan, S.; Haq, S.; Sharma, A.; Joshi, T. T.; Moazam, H.; Miller, H.; Zaharia, M.; and Potts, C. 2023. DSPy: Compiling Declarative Language Model Calls into Self-Improving Pipelines. *arXiv preprint arXiv:2310.03714*.
- Kočíšký, T.; Schwarz, J.; Blunsom, P.; Dyer, C.; Hermann, K. M.; Melis, G.; and Grefenstette, E. 2018. The narrativeqa reading comprehension challenge. *Transactions of the Association for Computational Linguistics*, 6: 317–328.
- Kwiatkowski, T.; Palomaki, J.; Redfield, O.; Collins, M.; Parikh, A.; Alberti, C.; Epstein, D.; Polosukhin, I.; Devlin, J.; Lee, K.; Toutanova, K.; Jones, L.; Kelcey, M.; Chang, M.-W.; Dai, A. M.; Uszkoreit, J.; Le, Q.; and Petrov, S. 2019. Natural Questions: A Benchmark for Question Answering Research. *Transactions of the Association for Computational Linguistics*, 7: 452–466.
- Li, X.; Sun, X.; Meng, Y.; Liang, J.; Wu, F.; and Li, J. 2020. Dice Loss for Data-imbalanced NLP Tasks. In Jurafsky, D.; Chai, J.; Schluter, N.; and Tetreault, J., eds., *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 465–476. Online: Association for Computational Linguistics.
- Li, Y.; Su, H.; Shen, X.; Li, W.; Cao, Z.; and Niu, S. 2017. DailyDialog: A Manually Labelled Multi-turn Dialogue Dataset. In *Proceedings of the Eighth International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, 986–995.
- Lin, C.-Y. 2004. ROUGE: A Package for Automatic Evaluation of Summaries. In *Text Summarization Branches Out*, 74–81. Barcelona, Spain: Association for Computational Linguistics.
- Liu, J.; Shen, D.; Zhang, Y.; Dolan, B.; Carin, L.; and Chen, W. 2022. What Makes Good In-Context Examples for GPT-3? In Agirre, E.; Apidianaki, M.; and Vulić, I., eds., *Proceedings of Deep Learning Inside Out (DeeLIO 2022): The 3rd Workshop on Knowledge*

- Extraction and Integration for Deep Learning Architectures*, 100–114. Dublin, Ireland and Online: Association for Computational Linguistics.
- Liu, Y.; Fabbri, A.; Liu, P.; Zhao, Y.; Nan, L.; Han, R.; Han, S.; Joty, S.; Wu, C.-S.; Xiong, C.; and Radev, D. 2023. Revisiting the Gold Standard: Grounding Summarization Evaluation with Robust Human Evaluation. In Rogers, A.; Boyd-Graber, J.; and Okazaki, N., eds., *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 4140–4170. Toronto, Canada: Association for Computational Linguistics.
- Madaan, A.; Tandon, N.; Gupta, P.; Hallinan, S.; Gao, L.; Wiegrefe, S.; Alon, U.; Dziri, N.; Prabhunoye, S.; Yang, Y.; et al. 2024. Self-refine: Iterative refinement with self-feedback. *Advances in Neural Information Processing Systems*, 36.
- MetaAI. 2024. LLaMA 3 Model. Accessed: 2024-06-13.
- Mu, W.; and Lim, K. H. 2022. Universal Evasion Attacks on Summarization Scoring. In Bastings, J.; Belinkov, Y.; Elazar, Y.; Hupkes, D.; Saphra, N.; and Wiegrefe, S., eds., *Proceedings of the Fifth BlackboxNLP Workshop on Analyzing and Interpreting Neural Networks for NLP*, 104–118. Abu Dhabi, United Arab Emirates (Hybrid): Association for Computational Linguistics.
- Nishida, K.; Saito, I.; Nishida, K.; Shinoda, K.; Otsuka, A.; Asano, H.; and Tomita, J. 2019. Multi-style Generative Reading Comprehension. In Korhonen, A.; Traum, D.; and Márquez, L., eds., *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 2273–2284. Florence, Italy: Association for Computational Linguistics.
- Nori, H.; King, N.; McKinney, S. M.; Carignan, D.; and Horvitz, E. 2023. Capabilities of gpt-4 on medical challenge problems. *arXiv preprint arXiv:2303.13375*.
- Pal, A.; Umapathi, L. K.; and Sankarasubbu, M. 2022. Medmcqa: A large-scale multi-subject multi-choice dataset for medical domain question answering. In *Conference on health, inference, and learning*, 248–260. PMLR.
- Pan, L.; Saxon, M.; Xu, W.; Nathani, D.; Wang, X.; and Wang, W. Y. 2023. Automatically correcting large language models: Surveying the landscape of diverse self-correction strategies. *arXiv preprint arXiv:2308.03188*.
- Prasad, A.; Hase, P.; Zhou, X.; and Bansal, M. 2023. GrIPS: Gradient-free, Edit-based Instruction Search for Prompting Large Language Models. In Vlachos, A.; and Augenstein, I., eds., *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, 3845–3864. Dubrovnik, Croatia: Association for Computational Linguistics.
- Pryzant, R.; Iter, D.; Li, J.; Lee, Y.; Zhu, C.; and Zeng, M. 2023. Automatic Prompt Optimization with “Gradient Descent” and Beam Search. In Bouamor, H.; Pino, J.; and Bali, K., eds., *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, 7957–7968. Singapore: Association for Computational Linguistics.
- Rajpurkar, P.; Zhang, J.; Lopyrev, K.; and Liang, P. 2016. SQuAD: 100,000+ Questions for Machine Comprehension of Text. In Su, J.; Duh, K.; and Carreras, X., eds., *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, 2383–2392. Austin, Texas: Association for Computational Linguistics.
- Reimers, N.; and Gurevych, I. 2019. Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics.
- Shinn, N.; Cassano, F.; Gopinath, A.; Narasimhan, K.; and Yao, S. 2023. Reflexion: language agents with verbal reinforcement learning. In Oh, A.; Naumann, T.; Globerson, A.; Saenko, K.; Hardt, M.; and Levine, S., eds., *Advances in Neural Information Processing Systems*, volume 36, 8634–8652. Curran Associates, Inc.
- Sordoni, A.; Yuan, E.; Côté, M.-A.; Pereira, M.; Trischler, A.; Xiao, Z.; Hosseini, A.; Níedtnér, F.; and Le Roux, N. 2024. Joint prompt optimization of stacked llms using variational inference. *Advances in Neural Information Processing Systems*, 36.
- Touvron, H.; Martin, L.; Stone, K.; Albert, P.; Almahairi, A.; Babaei, Y.; Bashlykov, N.; Batra, S.; Bhargava, P.; Bhosale, S.; et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- Van der Maaten, L.; and Hinton, G. 2008. Visualizing data using t-SNE. *Journal of machine learning research*, 9(11).
- Wang, B.; Liu, Z.; and Chen, N. 2023. Instructive Dialogue Summarization with Query Aggregations. In Bouamor, H.; Pino, J.; and Bali, K., eds., *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, 7630–7653. Singapore: Association for Computational Linguistics.
- Wang, X.; Li, C.; Wang, Z.; Bai, F.; Luo, H.; Zhang, J.; Jojic, N.; Xing, E.; and Hu, Z. 2023. PromptAgent: Strategic Planning with Language Models Enables Expert-level Prompt Optimization. In *The Twelfth International Conference on Learning Representations*.
- Wei, J.; Wang, X.; Schuurmans, D.; Bosma, M.; Xia, F.; Chi, E.; Le, Q. V.; Zhou, D.; et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35: 24824–24837.
- Xu, H.; Chen, Y.; Du, Y.; Shao, N.; Yanggang, W.; Li, H.; and Yang, Z. 2022. GPS: Genetic Prompt Search for Efficient Few-Shot Learning. In Goldberg, Y.; Kozareva, Z.; and Zhang, Y., eds., *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, 8162–8171. Abu Dhabi, United Arab Emirates: Association for Computational Linguistics.
- Yan, J. N.; Liu, T.; Chiu, J.; Shen, J.; Qin, Z.; Yu, Y.; Lakshmanan, C.; Kurzion, Y.; Rush, A.; Liu, J.; and Bendersky, M. 2024. Predicting Text Preference Via Structured Comparative Reasoning. In Ku, L.-W.; Martins, A.; and Srikumar, V., eds., *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 10040–10060. Bangkok, Thailand: Association for Computational Linguistics.
- Yang, C.; Wang, X.; Lu, Y.; Liu, H.; Le, Q. V.; Zhou, D.; and Chen, X. 2023. Large Language Models as Optimizers. In *The Twelfth International Conference on Learning Representations*.
- Yim, W.-w.; Fu, Y.; Ben Abacha, A.; Snider, N.; Lin, T.; and Yetisgen, M. 2023. Aci-bench: a novel ambient clinical intelligence dataset for benchmarking automatic visit note generation. *Scientific Data*, 10(1): 586.
- Yuksekonul, M.; Bianchi, F.; Boen, J.; Liu, S.; Huang, Z.; Guestrin, C.; and Zou, J. 2024. TextGrad: Automatic “Differentiation” via Text.
- Zha, Y.; Yang, Y.; Li, R.; and Hu, Z. 2023. AlignScore: Evaluating Factual Consistency with A Unified Alignment Function. In Rogers, A.; Boyd-Graber, J.; and Okazaki, N., eds., *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 11328–11348. Toronto, Canada: Association for Computational Linguistics.
- Zhang, T.; Kishore, V.; Wu, F.; Weinberger, K. Q.; and Artzi, Y. 2019. BERTScore: Evaluating Text Generation with BERT. In *International Conference on Learning Representations*.
- Zhou, Y.; Muresanu, A. I.; Han, Z.; Paster, K.; Pitis, S.; Chan, H.; and Ba, J. 2022. Large Language Models are Human-Level Prompt Engineers. In *The Eleventh International Conference on Learning Representations*.