

# Dual-Phase Accelerated Prompt Optimization

Anonymous ACL submission

## Abstract

Gradient-free prompt optimization methods have made significant strides in enhancing the performance of closed-source Large Language Model (LLMs) across a wide range of tasks. However, existing approaches make light of the importance of high-quality prompt initialization and the identification of effective optimization directions, thus resulting in substantial optimization steps to obtain satisfactory performance. In this light, we aim to accelerate prompt optimization process to tackle the challenge of low convergence rate. We propose a dual-phase approach which starts with generating high-quality initial prompts by adopting a well-designed meta-instruction to delve into task-specific information, and iteratively optimize the prompts at the sentence level, leveraging previous tuning experience to expand prompt candidates and accept effective ones. Extensive experiments on eight datasets demonstrate the effectiveness of our proposed method, achieving a consistent accuracy gain over baselines with less than five optimization steps.

## 1 Introduction

LLMs have demonstrated remarkable capabilities across a wide range of natural language processing (NLP) tasks, including machine translation (Qin et al., 2024), summarization (Goyal et al., 2022), and question answering (Zhang et al., 2023a). The dependency on prompt quality has led to the emergence of prompt engineering (Diao et al., 2023b; White et al., 2023), aiming at crafting effective prompts to elicit the desired responses from LLMs. As the need for efficient prompt design becomes increasingly evident (Liu et al., 2021b), automatic prompt optimization has been introduced to streamline the prompt design process, ensuring that LLMs are utilized to their full potential (Gao et al., 2021; Liu et al., 2021a; Reynolds and McDonnell, 2021).

Automatic prompt optimization can be broadly categorized into gradient-based and gradient-free




Figure 1: Average accuracy improvement on eight datasets with *four* optimization steps.

methods. Gradient-based methods (Shin et al., 2020; Li and Liang, 2021; Liu et al., 2021b, 2022) are devised for open-source LLMs to enable the optimization of prompts through adjustments based on model gradient. Gradient-free methods have emerged as the predominant approach for closed-source LLMs, which focuses on refining prompts without access to the model gradient (Prasad et al., 2022; Yang et al., 2023b; Guo et al., 2023). Starting from initial prompts, these methods usually expand candidate prompts using searching methods (?) and then accepting the more prominent ones in an iterative manner. This paper focuses on gradient-free methods due to the distinguished abilities of closed-source LLMs and the challenge of optimizing their prompts with limited model information.

We argue that current gradient-free prompt optimization methods have not adequately considered the rate of convergence. Typically, these methods demand an excessive number of optimization steps to obtain satisfactory prompts due to the limited access to model details, the vast discrete search space, and the uncertain optimization directions (Wang et al., 2023; Pan et al., 2023; Yang et al., 2023b). Representative work such as OPRO (Yang et al., 2023b) even necessitates nearly 200 optimization steps for some NLP tasks. This requirement for excessive optimization steps makes existing methods impractical for real-world applications since users are understandably reluctant to tolerate extensive optimization steps to achieve satisfactory performance levels. Therefore, we aim to achieve accel-

erated prompt optimization, obtaining satisfactory performance via few optimization steps (*e.g.*,  $< 5$ ).

To achieve accelerated prompt optimization, two crucial factors need to be considered: high-quality initial prompts and effective optimization directions. Firstly, the initialization of the prompt plays a crucial role in determining the efficiency of the optimization process (Ye et al., 2023), whereas existing approaches pay insufficient attention to the impact of initialization on subsequent optimization steps. Therefore, we aim to obtain initial prompts of high quality, laying a solid foundation for the accelerated optimization process. Secondly, the accelerated prompt optimization needs to identify the most effective optimization directions in each step, streamlining efficient optimization from the initial prompts. Thus, we aim to design a more refined expansion tuned by experience and acceptance of candidate prompts enhanced by examination of past failure cases.

To this end, we propose a dual-phase approach to achieve the accelerated gradient-free prompt optimization. Our approach consists of two phases: high-quality initial prompt generation, and experience-tuned optimization. Firstly, we utilize a well-designed meta-instruction to guide the LLM in generating high-quality and structured initial prompts that contain task-specific information, including task type and description, output format and constraints, suggested reasoning process, and professional tips. After that, we devise a sentence-level prompt optimization strategy for efficiently optimization on the long initial prompt, leveraging previous direction tuning experience, together with failure cases, to select sentences in the initial prompt to be expanded and accept effective prompt candidates. Extensive experiments on two different LLMs across eight datasets confirm the effectiveness and superiority of our method. Our contributions are threefold:

- We reveal the issue of low convergence rate in gradient-free prompt optimization, and highlight the problem of accelerated prompt optimization.
- We propose a dual-phase approach, achieving accelerated prompt optimization through high-quality initial prompt generation and experience-tuned optimization.
- We conduct extensive experiments, demonstrating that the proposed method achieves satisfying performance within few optimization steps.

## 2 Related Work

The gradient-free prompt optimization for closed-source LLMs typically contains two phases: initialization and iterative optimization steps, where the optimization step consists of expansion and selection stages.

**Initialization.** The prompt initialization for optimization can be achieved manually or autonomously. Manual initialization often entails professional machine learning engineers formulating prompts, as delineated in (Pryzant et al., 2023). Concurrently, works such as (Guo et al., 2023), (Pan et al., 2023), and (Wang et al., 2023) utilize existing manual prompts as the foundational set to harness human creativity. In contrast, automated initialization leverages the power of LLM generation, which is exemplified by (Zhang et al., 2023b), generating prompts from few-shot exemplars and a rudimentary description, and (Zhou et al., 2022), fabricating prompts based on meta-prompts and illustrative input-output examples. Our method belongs to the automated initialization, improving the initial prompt generation for acceleration.

**Optimization.** The optimization step is achieved by expanding prompt candidates through modifying from the initial prompt, and selecting the better candidates for the next iteration. The expansion stage can be executed through rephrasing, as in (Zhou et al., 2022), where high-scoring prompts undergo evolution akin to a Monte Carlo search methodology, or through heuristic algorithms that automatically revise prompts, as in (Guo et al., 2023) and (Pan et al., 2023). More complex regeneration strategies are employed by works like (Wang et al., 2023), where the optimizer LLM progressively expands prompts based on task delineations and historical iterations. The expansion can also be implemented leveraging an open-source LLM (Lin et al., 2023; Chen et al., 2024). Reinforcement learning-based methods have also been adopted for prompt modification (Diao et al., 2023a). Moreover, the granularity of prompt modification exhibits variation across studies. Heuristic-based methods and (Hsieh et al., 2023) work operate at the word/token granularity, while classical optimization algorithms like (Pryzant et al., 2023; Zhou et al., 2022) consider the entire prompt. The selection stage generally utilized the performance of the prompt on a held-out validation set (Pryzant et al., 2023; Zhou et al., 2022; Wang et al., 2023),

while recent work also explores human preference feedback (Lin et al., 2024) or score feedback from other LLMs (Yang et al., 2024).

### 3 Problem formulation

#### 3.1 Gradient-Free Prompt Optimization

For a target NLP task  $\mathcal{T}$  with input  $x$ , the closed-source LLM predicts the output  $\hat{y}$  given  $x$  concatenated with the prompt  $p$ , where  $x$ ,  $\hat{y}$  and  $p$  are all word sequences. The aim for prompt optimization is to find an optimal prompt  $p^*$  that obtains the desired  $\hat{y}$ , which can be evaluated by metrics such as accuracy with reference to the ground truth  $y$ . The gradient-free prompt optimization contains an initialization phase followed by  $K$  iterative optimization steps. The  $k$ -th optimization step starts from an initial prompt  $p_{k-1}$ ,  $k \in [1, K]$ , and sequentially performs two stages: expansion of prompt candidates, and acceptance of the prominent prompts as the next initial prompts, as detailed below.

**Expansion of Prompt Candidates.** At the  $k$ -th optimization step, The expansion stage search for new prompt candidates with potential better performance starting from  $p_{k-1}$ , with searching methods such as edit-based (Prasad et al., 2022) and LLM rewriting (Pryzant et al., 2023). Formally, the expansion function  $f_E(\cdot)$  generates prompt candidate set  $P_k^c = \{p_{k_1}^c, \dots, p_{k_Q}^c\}$  with size  $Q$ .

$$P_k^c = f_E(p_{k-1}). \quad (1)$$

**Acceptance of Prominent Prompts.** The acceptance stage evaluates the performance of each prompt candidate in  $P_k^c$  to determine whether it should be continued for next optimization step. This is usually achieved by evaluation on a held-out validation set  $V = \{(x^v, y^v)\}$ , and accepting the top-performing prompt candidates. Formally, with the evaluation function on LLM as  $f_S(\cdot)$ ,

$$\begin{aligned} r_i^k &= f_S(p_{k_i}^c, V), i \in [1, \dots, Q], \\ p_k &= p_{k_j}^c, \text{ where } j = \operatorname{argmax}(\{r_1^k, \dots, r_Q^k\}). \end{aligned} \quad (2)$$

where  $\operatorname{argmax}(\cdot)$  denotes the index of the maximum value. At the final optimization steps, the top-performing prompt  $p_K$  will be accepted as the optimized prompt  $p^*$ .

#### 3.2 Accelerated Prompt Optimization

Although current research on gradient-free prompt optimization can achieve significant performance

gains on multiple tasks, demands for a great number of optimization steps hinder their practicability in real-world scenarios. For instance, Yang et al. (2023b) does not converge even after over 150 steps in some tasks; Wang et al. (2023) finds a good solution in 50 to 75 steps. Therefore, we highlight the problem of accelerated prompt optimization, *i.e.*, obtaining  $p^*$  with satisfactory performance in few optimization steps, *e.g.*,  $K < 5$ .

## 4 Proposed method

### 4.1 Motivation

We believe that two factors are crucial for achieving accelerated prompt optimization, which current gradient-free prompt optimization methods fail to achieve. Firstly, the initial prompt  $p_0$  plays a crucial role in accelerating the prompt optimization process (Ye et al., 2023), where  $p_0$  with better LLM performance makes the optimization towards better prompts easier, preventing LLMs from excessively exploring suboptimal prompt regions. This is generally overlooked by existing research that utilizes uninformative initial prompts, *e.g.*, (Yang et al., 2023b). Therefore, we propose to devise high-quality  $p_0$  by crafting a novel initial prompt schema. Furthermore, a more precise expansion and acceptance of prompt candidates ensure highly efficient optimization direction and fewer optimization steps. Current expansion and acceptance techniques optimize the prompt towards improving the general task performance, where effective optimization direction in each step is hard to ensure. To tackle this, we propose to utilize the past failure cases from previous optimization steps to further navigate the expansion and acceptance of prompt candidates. We illustrate our dual-phase approach as follows (*cf.* Figure 2).

### 4.2 High-Quality Initial Prompt Generation

We think that a high-quality initial prompt that can elicit the desired response from LLMs should be able to provide clear task instruction and detailed task-related information. Specifically, it should 1) give a clear definition of the task type and provide a detailed task description, 2) define the output format and constraints, 3) provide insights on the reasoning processes and professional tips. To achieve such initial prompts, we are inspired by the step-back prompting (Zheng et al., 2023) which demonstrates LLM’s ability in deriving high-level concepts and principles from examples. Thus,

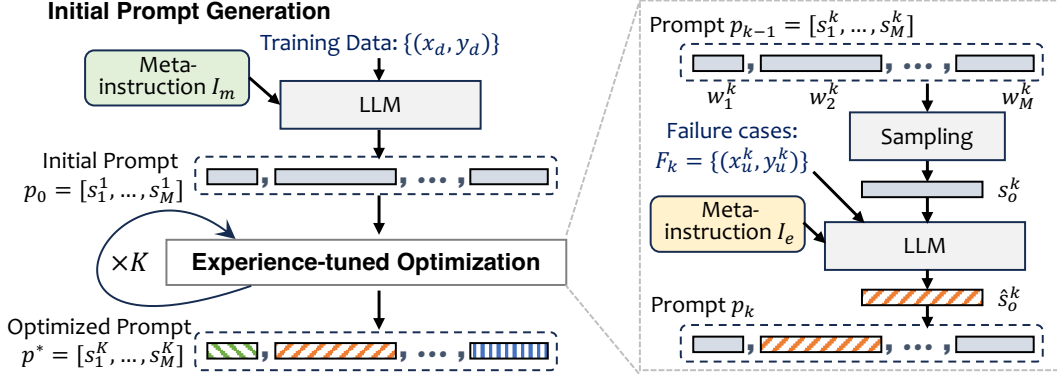


Figure 2: Illustration of the proposed method.

following (Zhou et al., 2022), we design a meta-instruction  $I_m$  (cf. Figure 3), leveraging LLM’s ability to generate  $p_0$  by observing the input-output exemplars of the target task  $\mathcal{T}$  and inferring the above required information. Formally, defining input-output exemplars as  $D = \{(x_d, y_d)\}$ ,

$$p_0 = LLM(I_m, D). \quad (3)$$

### 4.3 Experience-tuned Optimization

In the optimization phase, it is necessary to tune the expansion and acceptance of prompt candidates to quickly improve the task performance as evaluated on the validation set  $V$  and thus reduce optimization steps. Inspired by previous research (Pryzant et al., 2023), we intend to make the best of past failure cases to generate promising prompt candidates and filter out unnecessary optimization attempts. In each optimization steps, we maintain a failure case set  $F_k = \{(x_k^f, y_k^f)\}$  containing the examples from  $V$  where the initial prompt  $p_{k-1}$  fails to predict the ground-truth in the acceptance stage, i.e.,  $\hat{y}_k^f \neq y_k^f$ .

**Expansion.** In the expansion stage, since the initial prompts are long prompts with at least four sentences, we aim to improve the expansion efficiency by segmenting them into individual sentences for sentence-level expansion following LongPO (Hsieh et al., 2023). Moreover, since different sentences in the initial prompts contains different task-related information and may have different impact on the task performance, we devise sentence weights  $w^k$  to estimate the impact of each sentence on the performance improvement, which is updated leveraging the past failure cases. We first split the initial prompt  $p_0$  into  $M$  sentences, and initialize the

weight  $w^1$  for each sentence as 1.

$$p_0 = [s_1^1, s_2^1, \dots, s_M^1], \quad (4)$$

$$w_t^1 = 1, t \in [1, M].$$

In the  $k$ -th optimization step, we compute the acceptance probability  $\text{Pr}^k$  for each sentence:

$$\text{Pr}_i^k = \frac{\exp(w_i^k)}{\sum_{j=1}^M \exp(w_j^k)}. \quad (5)$$

After that, we sample a sentence for expansion based on the probability distribution  $\text{Pr}^k = [\text{Pr}_1^k, \dots, \text{Pr}_M^k]$ , where the sampled sentence is denoted as  $s_o^k, o \in [1, M]$ . For expansion of  $s_o^k$ , we design a meta-instruction  $I_e$  (cf. Figure 4) to instruct LLM to generate a revised sentence considering the past experience.

$$\hat{s}_o^k = LLM(I_e, p_{k-1}, F_k, s_o^k). \quad (6)$$

Before passing  $\hat{s}_o^k$  to the acceptance stage, we design additional strategies to further guarantee the effectiveness of the generated sentence leveraging  $F_k$ . Firstly, to ensure  $\hat{s}_o^k$  can actually improve over  $s_o^k$ , we replace  $s_o^k$  in  $p_{k-1}$  with  $\hat{s}_o^k$ , denoted as  $\hat{p}_k$ , and evaluate whether  $\hat{p}_k$  outperforms  $p_{k-1}$  on  $F_k$ . We accept  $\hat{s}_o^k$  only when  $\hat{p}_k$  has improved the performance over  $p_{k-1}$  larger than a threshold  $H_F$ .

$$f_S(\hat{p}_k, F_k) - f_S(p_{k-1}, F_k) > H_F. \quad (7)$$

Besides, to avoid repeatedly generating the same ineffective  $\hat{s}_o^k$ , we build a collection  $\mathcal{G}$  of undesired sentence revisions and check whether  $\hat{s}_o^k$  has appeared in  $\mathcal{G}$ . If the above two criteria are not met, we abandon  $\hat{s}_o^k$  and regenerate starting from Eq. 6.

**Acceptance.** In addition to evaluating  $\hat{p}_k$ ’s performance on the entire failure case  $F_k$ , we also

#### meta-instruction for initialization

You gave me an instruction on a certain task and some example inputs with chain-of-thought. I read the instruction carefully and wrote an output with chain-of-thought for every input correctly. Here are some correct input-output pairs which strictly meet all your requirements:

{example\_pairs}

The instruction given contains the following parts. Based on the input-output pairs provided, give me the final complete instruction in English without any explanation:

###Task type###

Task type: This is a <...> task.

###Task detailed description###

Task detailed description: <Task detailed description>

###Your output must satisfy the following format and constraints###

Output format(type): <Output format or its type>

Output constraints: <constraints on output>

###You must follow the reasoning process###

<add several reasoning steps if it's necessary>

###Tips###

<add several useful tips from a professional point of view to accomplish this task better>

Figure 3: Meta-instruction used in our initialization phase to generate high-quality initial prompts.

331 evaluate its performance on the validation set  $V$ .  
332 We accept  $\hat{p}_k$  as the next initial prompt  $p_k$  only  
333 when  $\hat{p}_k$  has improved the performance over  $p_{k-1}$   
334 larger than a threshold  $H_V$ . Otherwise, we abandon  
335  $\hat{p}^k$  and restart from sampling  $s_o^k$ .

$$336 \quad f_S(\hat{p}_k, V) - f_S(p_{k-1}, V) > H_V. \quad (8)$$

337 If  $\hat{p}^k$  is accepted, we update its sentence weights.  
338 We calculate the mixed evaluation result  $f_R(\cdot)$  and  
339 update the  $w^{k+1}$  as follows, where  $\alpha$  and the learn-  
340 ing rate  $\eta$  are adjusting hyperparameters.

$$341 \quad f_R(\hat{p}_k) = \alpha f_S(\hat{p}_k, V) + (1 - \alpha) f_S(\hat{p}_k, F_k). \quad (9)$$

$$342 \quad w_i^{k+1} = w_i^k \exp\left(\frac{\eta f_R(\hat{p}_k)}{\text{Pr}_i^k M}\right).$$

343 When the number of times that Eq. 7 or Eq. 8 is  
344 not satisfied accumulates to 5, we consider the al-  
345 gorithm to have converged.

#### meta-instruction for optimization

I'm trying to write a zero-shot prompt which consists of four parts.

My current prompt is:

[{prompt\_to\_revise}]

But it gets the following outputs that fail to match the expected outputs:

{failed\_cases}

The sentence I want to revise is:

{sentences[chosen\_sentence]}

Comparing the wrong outputs with their corresponding expected answers under the same input, optimize the above sentence to help AI understand the task more comprehensively and accomplish this task better.

Your response format is as follows.

The given sentence

'{sentences[chosen\_sentence]}' should be revised as:

Figure 4: Meta-instruction used in the optimization phase.

The weight formula is designed to adaptively update the importance of each sentence in the prompt based on its impact on overall performance improvement.  $f_R(\hat{p}_k)$  modulates the magnitude of the weight adjustment: a higher  $f_R(\hat{p}_k)$  leads to larger updates.  $\text{Pr}_i^k$  determines the weight's contribution, while  $M$  is used for normalization to ensure balanced weight updates. The learning rate  $\eta$  controls the extent of weight adjustments based on the evaluation feedback. Inspired by the EXP3 algorithm, these components facilitate a dynamic and adaptive optimization process, tuned by empirical performance data. The whole process is summarized in Algorithm 1.

## 5 Experiments

In this section, we begin by detailing datasets, baselines, and the implementation of the experiments. Following this, we conduct comprehensive and controlled experiments on our method.

### 5.1 Experimental Settings

**Datasets.** Our experiments are first conducted on general natural language understanding tasks across four datasets to validate our method, specifically focusing on sentiment classification (SST-2 (Socher et al., 2013)), topic classification (AG's News (Zhang et al., 2015)), TREC (Voorhees and Tice, 2000) and subjectivity classification (Subj (Pang and Lee, 2004)). Then we perform our ap-

---

**Algorithm 1**Dual-Phase Accelerated Prompt Optimization

---

**Require:** Input-output exemplars  $D$ , validation set  $V$ , meta-instruction  $I_m$  and  $I_e$ .**Ensure:** Optimized prompt  $p^*$ 

- 1: Initialize  $p_0$  (Eq. 3), derive failure case set  $F_1$
  - 2: Split  $p_0$  into  $M$  sentences  $[s_1^1, s_2^1, \dots, s_M^1]$ , initialize sentence weights  $\{w_i^1\}_{i=1}^M \leftarrow 1, k \leftarrow 1$
  - 3: **while** not converged **do**
  - 4:    $\triangleright$  **Expansion**
  - 5:   Sample a sentence  $s_o^k$  based on  $\text{Pr}^k$  (Eq. 5)
  - 6:   Generate revised sentence  $\hat{s}_o^k$  (Eq. 6)
  - 7:   Replace  $s_o^k$  in  $p_{k-1}$  with  $\hat{s}_o^k$  to get  $\hat{p}_k$
  - 8:   **if**  $\hat{s}_o^k \in \mathcal{G}$  **or** (Eq. 7) is not satisfied **then**
  - 9:     Add  $\hat{s}_o^k$  to  $\mathcal{G}$
  - 10:     Regenerate  $\hat{s}_o^k$  from line 6
  - 11:   **end if**
  - 12:    $\triangleright$  **Acceptance**
  - 13:   **if** (Eq. 8) is not satisfied **then**
  - 14:     Restart from line 5
  - 15:   **end if**
  - 16:    $p_k \leftarrow \hat{p}_k$ , update  $w_i^{k+1}, k \leftarrow k + 1$
  - 17:   Update  $F_k$  with new failure cases
  - 18: **end while**
  - 19: **return** optimized prompt  $p^* = p_k$
- 

proach to the challenging BBH tasks (Suzgun et al., 2022), which include manually provided few-shot Chain-of-Thought (CoT) prompts containing task descriptions and demonstrations.

**Baselines.** We compare our method with three popular prompt optimization methods for zero-shot black-box prompting and the well-crafted prompts manually provided in BBH tasks: **APO** (Pryzant et al., 2023): Generating natural language “gradients” to criticize and improve the current prompts. **APE** (Zhou et al., 2022): Proposing both a naive and an iterative Monte Carlo search methods to approximate the solution to the prompt optimization problem. **PromptAgent** (Wang et al., 2023): Automating expert-level prompt generation by treating it as a strategic planning problem using Monte Carlo tree search and error feedback to refine and optimize prompts. **Manual Prompt** (Suzgun et al., 2022): The few-shot CoT version of human-designed prompts with teaching examples developed in BBH tasks.

**Implementation details.** In line with (Wang et al., 2023), since BBH tasks lack an official train-test split, we shuffle the data and allocate approxi-

mately half for testing. The rest is used for training, prompt generation, and optimization. For datasets with predefined test sets, we use those directly.

Unless otherwise stated, we evaluate performance (*i.e.*, testing accuracy) on GPT-3.5-Turbo using the OpenAI API<sup>1</sup> (currently gpt-3.5-turbo-0125) in a zero-shot prompt setting. The temperature is set to 0 for prediction and 0.5 for prompt generation to enhance diversity. To accelerate prompt optimization, we limit the maximum optimization steps to **four** for all methods, while keeping other baseline parameters and settings at default. At the beginning of prompt initialization, eight exemplars are obtained by concatenating unique input-output pairs from the shuffled training data until the desired amount is reached, ensuring no duplicate inputs. Due to limited computational resources, our approach generates and optimizes only one initial prompt. By default, we set  $H_F = 0.3$ ,  $H_V = 0.1$ ,  $\alpha = 0.4$ , and  $\eta = 0.055$  in Algorithm 1 to accelerate the optimization phase.

## 5.2 Main Results & Analysis

Task	Few-shot		Zero-shot		
	Manual	APO	APE	PA	Ours
SST-2	/	0.89	<u>0.92</u>	0.443	<b>0.978</b>
AG’s News	/	0.88	0.819	0.785	<b>0.928</b>
TREC	/	<b>0.795</b>	0.513	0.687	<u>0.785</u>
Subj	/	<u>0.64</u>	0.593	0.494	<b>0.72</b>
Logical Five	0.388	0.392	0.404	<u>0.443</u>	<b>0.48</b>
Hyperbaton	0.744	0.808	<u>0.865</u>	0.823	<b>0.88</b>
Disambiguation	0.580	0.688	0.645	<u>0.696</u>	<b>0.74</b>
Salient Translation	0.544	0.456	<u>0.538</u>	0.468	<b>0.548</b>
Avg.	0.564	0.694	0.662	0.605	<b>0.757</b>

Table 1: Accuracy on eight tasks on GPT-3.5-Turbo. PA indicates PromptAgent. Bold and underlined text indicate the best and second-best results, respectively.

**Overall Results.** Table 1 demonstrates the effectiveness of our accelerated dual-phase approach across 8 NLP tasks compared to classic prompt optimization methods. Our method significantly outperforms all baselines, achieving an average improvement of approximately **10.7%** over APO, **16.4%** over APE, and **29.7%** over PromptAgent across the given tasks.

Our method also surpasses few-shot CoT human-crafted prompts with an approximately **17.6%** average improvement on selected BBH tasks, indicating its ability to produce high-quality prompts that enhance the black-box LLM’s capabilities in logical

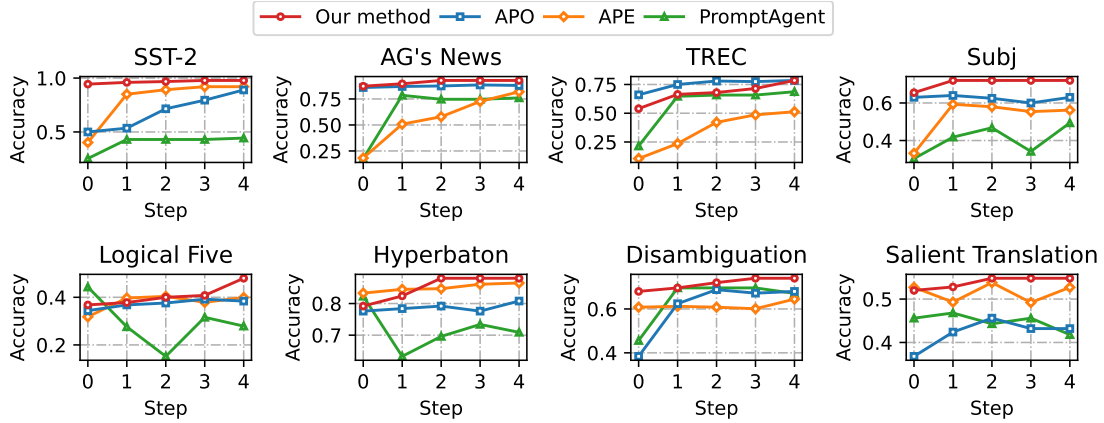


Figure 5: Performance (accuracy) over 4 steps across 8 tasks on GPT-3.5-Turbo.

deduction, grammar, language understanding, and multilingual tasks without teaching examples.

**Analysis.** To understand this result, we analyzed the prompt expansion and acceptance processes: In prompt expansion, our method leverages past experience, filters out unnecessary optimization attempts, and collects undesired revisions. This contrasts with baseline methods that inefficiently explore prompt space and underutilize past iterations. APE lacks reflection on past iterations, slowing its Monte Carlo-based search. APO uses error feedback to guide beam search but is slowed by evaluating many paths. PromptAgent’s Monte Carlo Search Tree explores prompt optimization through simulations, but limited steps lead to suboptimal results.

In the acceptance process, inspired by the EXP3 algorithm (Auer et al., 1995), our method uses weighted sentences and modifications to enhance prompt quality, making it superior in identifying promising candidates and optimizing directions.

**Convergence analysis.** To evaluate our method’s convergence within four steps compared to others, we examine how quickly each method achieves peak performance across datasets. Figure 5 shows the performance (accuracy) variation of four prompt optimization methods across eight datasets, with each subfigure representing a different dataset. While APO, APE, and PromptAgent experience fluctuations or plateau at lower accuracy, our method demonstrates the fastest convergence across most datasets, often reaching near-peak performance within the first two steps. This rapid convergence highlights our method’s efficiency in optimizing prompts quickly and effectively, making it promising for tasks requiring




Figure 6: Results on GPT-3.5-Turbo with different initial prompt schemas.




Figure 7: Results on GPT-3.5-Turbo with different optimization learning rates.

prompt optimization within a few steps.

### 5.3 Ablation Study

We conduct four ablation experiments to assess the efficacy of our method.

#### 5.3.1 Different Initial Prompt Schemas

Our method uses a meta-instruction to generate a prompt with four components: a) task type and description, b) output format and constraints, c) suggested reasoning process, and d) professional tips. We define: *Schema 4*: All four components *Schema 3*: First three components *Schema 2*: First two components *Schema 1*: Task type and description only (common in current techniques) We vary the meta-instructions for these schemas and conduct four-step prompt optimization experiments on SST-2 and AG’s News to assess their impact on optimization.

As shown in Figure 6, initial prompts from

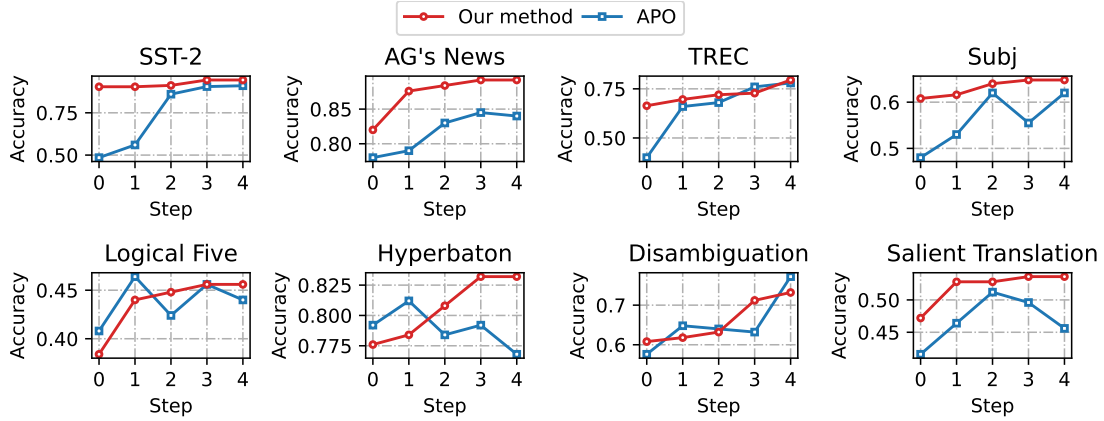


Figure 8: Accuracy over 4 steps across 8 tasks on Baichuan2-Turbo.

Schema 4 yield the highest evaluation results. In contrast, Schema 1 has the lowest metrics and often falls into suboptimal local minima, a common issue with current methods. This comparison validates our meta-instruction design and underscores that a high-quality initial prompt is crucial for quickly identifying the optimal prompt.

### 5.3.2 Sensitivity to Learning Rate

During the optimization phase, the learning rate  $\eta$  controls the extent of sentence weight updates after each round. A higher  $\eta$  results in significant updates and responsiveness to recent performance changes, while a lower  $\eta$  promotes stability with gradual adjustments. This balance is crucial for navigating the trade-off between exploration and exploitation.

We conduct prompt optimization experiments on SST-2 and AG’s News within four steps, testing  $\eta$  values from 0.01 to 0.1. As shown in Figure 7,  $\eta = 0.055$  and  $\eta = 0.07$  are the most and second most effective in accelerating optimization.

### 5.3.3 Performance on Different LLM

As Table 1 indicates, APO is the best baseline method. Therefore, we compare our method with APO using Baichuan2 (Yang et al., 2023a), a Chinese large language model accessed via API. We conduct prompt optimization experiments on eight NLP datasets across four optimization steps on Baichuan2-Turbo. Figure 8 illustrates the performance variation of both methods across different datasets as optimization steps progress. APO fails to converge within four steps and shows greater performance volatility compared to GPT-3.5-Turbo. In contrast, our method demonstrates rapid convergence and strong optimization acceleration.




Figure 9: Accuracy over 20 steps on GPT-3.5-Turbo.

### 5.3.4 Results without Step Constraint

We report the results of prompt optimization with a maximum of 20 steps on two general NLU tasks. As shown in Figure 9, the strongest baseline, APO, converges on the SST-2 task with slightly lower accuracy than our method. However, on the AG’s News task, APO’s performance fluctuates significantly and lags behind our method. Thus, our method demonstrates superior performance and faster convergence compared to existing methods, even with fewer optimization steps.

## 6 Conclusion

In this paper, we addressed the issue of low convergence rates in gradient-free prompt optimization methods for LLMs. Our proposed dual-phase approach effectively accelerates prompt optimization by generating high-quality initial prompts and leveraging tuning experience to navigate the optimization process. Extensive experiments on two LLMs across eight datasets demonstrated the superiority of our method in achieving satisfactory performance within few optimization steps. Our approach not only enhances the efficiency of prompt optimization but also improves the overall performance of LLMs in various NLP tasks. Future work will focus on further refining the optimization strategies and exploring their applications in more diverse and complex scenarios.



## 550 Limitation

551 We acknowledge some limitations despite the  
552 promising results of our research that could pave  
553 the way for future studies:

554 1) Our experiments were limited to general NLP  
555 tasks and did not assess performance on special-  
556 ized tasks requiring domain knowledge. 2) Our  
557 method relies on labeled task data for prompt gen-  
558 eration and evaluation, raising concerns about its  
559 robustness in personalized or scenarios lacking la-  
560 beled data. 3) Our experiments were confined to  
561 GPT-3.5-Turbo and Baichuan2-Turbo, leaving the  
562 effectiveness of our method on other large language  
563 models to be validated in future studies.

564 Further study may be needed to address these  
565 limitations so as to improve the applicability and  
566 robustness of our approach in broader and more  
567 complex real-world applications.

## 568 References

569 Peter Auer, Nicolò Cesa-Bianchi, Yoav Freund, and  
570 Robert E. Schapire. 1995. [Gambling in a rigged  
571 casino: The adversarial multi-arm bandit problem.](#)  
572 *In IEEE Annual Symposium on Foundations of Com-  
573 puter Science.*

574 Lichang Chen, Jiuai Chen, Tom Goldstein, Heng  
575 Huang, and Tianyi Zhou. 2024. [Instructzero: Ef-  
576 ficient instruction optimization for black-box large  
577 language models.](#)

578 Shizhe Diao, Zhichao Huang, Ruijia Xu, Xuechun  
579 Li, LIN Yong, Xiao Zhou, and Tong Zhang. 2023a.  
580 [Black-box prompt learning for pre-trained language  
581 models.](#) *Transactions on Machine Learning Re-  
582 search.*

583 Shizhe Diao, Pengcheng Wang, Yong Lin, and Tong  
584 Zhang. 2023b. [Active prompting with chain-  
585 of-thought for large language models.](#) *ArXiv,*  
586 *abs/2302.12246.*

587 Tianyu Gao, Adam Fisch, and Danqi Chen. 2021.  
588 [Making pre-trained language models better few-shot  
589 learners.](#) *In Annual Meeting of the Association for  
590 Computational Linguistics.*

591 Tanya Goyal, Junyi Jessy Li, and Greg Durrett. 2022.  
592 [News summarization and evaluation in the era of  
593 gpt-3.](#) *ArXiv,* *abs/2209.12356.*

594 Qingyan Guo, Rui Wang, Junliang Guo, Bei Li, Kaitao  
595 Song, Xu Tan, Guoqing Liu, Jiang Bian, Yujia Yang,  
596 Tsinghua University, and Microsoft Research. 2023.  
597 [Connecting large language models with evolutionary  
598 algorithms yields powerful prompt optimizers.](#) *ArXiv,*  
599 *abs/2309.08532.*

600 Cho-Jui Hsieh, Si Si, Felix X. Yu, and Inderjit S. Dhillon.  
601 2023. [Automatic engineering of long prompts.](#)  
602 *ArXiv,* *abs/2311.10117.*

603 Xiang Lisa Li and Percy Liang. 2021. [Prefix-tuning:  
604 Optimizing continuous prompts for generation.](#) *Pro-  
605 ceedings of the 59th Annual Meeting of the Associa-  
606 tion for Computational Linguistics and the 11th Inter-  
607 national Joint Conference on Natural Language Pro-  
608 cessing (Volume 1: Long Papers),* *abs/2101.00190.*

609 Xiaoqiang Lin, Zhongxiang Dai, Arun Verma, See-  
610 Kiong Ng, Patrick Jaillet, and Bryan Kian Hsiang  
611 Low. 2024. [Prompt optimization with human feed-  
612 back.](#) *arXiv preprint arXiv:2405.17346.*

613 Xiaoqiang Lin, Zhaoxuan Wu, Zhongxiang Dai,  
614 Wenyang Hu, Yao Shu, See-Kiong Ng, Patrick Jaillet,  
615 and Bryan Kian Hsiang Low. 2023. [Use your instinct:  
616 Instruction optimization using neural bandits coupled  
617 with transformers.](#) *ArXiv,* *abs/2310.02905.*

618 Pengfei Liu, Weizhe Yuan, Jinlan Fu, Zhengbao Jiang,  
619 Hiroaki Hayashi, and Graham Neubig. 2021a. [Pre-  
620 train, prompt, and predict: A systematic survey of  
621 prompting methods in natural language processing.](#)  
622 *ACM Computing Surveys,* *55:1 – 35.*

623 Xiao Liu, Kaixuan Ji, Yicheng Fu, Weng Lam Tam,  
624 Zhengxiao Du, Zhilin Yang, and Jie Tang. 2022. [P-  
625 tuning: Prompt tuning can be comparable to fine-  
626 tuning across scales and tasks.](#) *In Annual Meeting of  
627 the Association for Computational Linguistics.*

628 Xiao Liu, Yanan Zheng, Zhengxiao Du, Ming Ding,  
629 Yujie Qian, Zhilin Yang, and Jie Tang. 2021b. [Gpt  
630 understands, too.](#) *ArXiv,* *abs/2103.10385.*

631 Rui Pan, Shuo Xing, Shizhe Diao, Xiang Liu, Kashun  
632 Shum, Jipeng Zhang, and Tong Zhang. 2023.  
633 [Plum: Prompt learning using metaheuristic.](#) *ArXiv,*  
634 *abs/2311.08364.*

635 Bo Pang and Lillian Lee. 2004. [A sentimental education:  
636 Sentiment analysis using subjectivity summarization  
637 based on minimum cuts.](#) *ArXiv,* *cs.CL/0409058.*

638 Archiki Prasad, Peter Hase, Xiang Zhou, and Mohit  
639 Bansal. 2022. [Grips: Gradient-free, edit-based in-  
640 struction search for prompting large language models.](#)  
641 *ArXiv,* *abs/2203.07281.*

642 Reid Pryzant, Dan Iter, Jerry Li, Yin Tat Lee, Chen-  
643 guang Zhu, and Michael Zeng. 2023. [Automatic  
644 prompt optimization with "gradient descent" and  
645 beam search.](#) *In Conference on Empirical Methods  
646 in Natural Language Processing.*

647 Libo Qin, Qiguang Chen, Xiachong Feng, Yang Wu,  
648 Yongheng Zhang, Yinghui Li, Min Li, Wanxiang  
649 Che, and Philip S. Yu. 2024. [Large language models  
650 meet nlp: A survey.](#)

651 Laria Reynolds and Kyle McDonell. 2021. [Prompt  
652 programming for large language models: Beyond the  
653 few-shot paradigm.](#) *Extended Abstracts of the 2021*

654		CHI Conference on Human Factors in Computing Systems.	Qinyuan Ye, Maxamed Axmed, Reid Pryzant, and Fereshte Khani. 2023. Prompt engineering a prompt engineer. <i>ArXiv</i> , abs/2311.05661.	711
655				712
656	Taylor Shin, Yasaman Razeghi, Robert L Logan IV, Eric Wallace, and Sameer Singh. 2020. Eliciting knowledge from language models using automatically generated prompts. <i>ArXiv</i> , abs/2010.15980.		Haopeng Zhang, Xiao Liu, and Jiawei Zhang. 2023a. Extractive summarization via chatgpt for faithful summary generation. In <i>Conference on Empirical Methods in Natural Language Processing</i> .	714
657				715
658				716
659				717
660	Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D. Manning, A. Ng, and Christopher Potts. 2013. Recursive deep models for semantic compositionality over a sentiment treebank. In <i>Conference on Empirical Methods in Natural Language Processing</i> .		Xiang Zhang, Junbo Jake Zhao, and Yann LeCun. 2015. Character-level convolutional networks for text classification. In <i>Neural Information Processing Systems</i> .	718
661				719
662				720
663				
664				
665				
666	Mirac Suzgun, Nathan Scales, Nathanael Scharli, Sebastian Gehrmann, Yi Tay, Hyung Won Chung, Aakanksha Chowdhery, Quoc V. Le, Ed Huai hsin Chi, Denny Zhou, and Jason Wei. 2022. Challenging big-bench tasks and whether chain-of-thought can solve them. In <i>Annual Meeting of the Association for Computational Linguistics</i> .		Zhihan Zhang, Shuo Wang, W. Yu, Yichong Xu, Dan Iter, Qingkai Zeng, Yang Liu, Chenguang Zhu, and Meng Jiang. 2023b. Auto-instruct: Automatic instruction generation and ranking for black-box language models. In <i>Conference on Empirical Methods in Natural Language Processing</i> .	721
667				722
668				723
669				724
670				725
671				726
672				
673	Ellen M. Voorhees and Dawn M. Tice. 2000. Building a question answering test collection. In <i>Annual International ACM SIGIR Conference on Research and Development in Information Retrieval</i> .		Huaixiu Steven Zheng, Swaroop Mishra, Xinyun Chen, Heng-Tze Cheng, Ed Huai hsin Chi, Quoc V. Le, and Denny Zhou. 2023. Take a step back: Evoking reasoning via abstraction in large language models. <i>ArXiv</i> , abs/2310.06117.	727
674				728
675				729
676				730
677	Xinyuan Wang, Chenxi Li, Zhen Wang, Fan Bai, Hao-tian Luo, Jiayou Zhang, Nebojsa Jojic, Eric P. Xing, and Zhiting Hu. 2023. Promptagent: Strategic planning with language models enables expert-level prompt optimization. <i>ArXiv</i> , abs/2310.16427.		Yongchao Zhou, Andrei Ioan Muresanu, Ziwen Han, Keiran Paster, Silviu Pitis, Harris Chan, and Jimmy Ba. 2022. Large language models are human-level prompt engineers. <i>ArXiv</i> , abs/2211.01910.	731
678				732
679				733
680				734
681				735
682	Jules White, Quchen Fu, Sam Hays, Michael Sandborn, Carlos Olea, Henry Gilbert, Ashraf Elnashar, Jesse Spencer-Smith, and Douglas C. Schmidt. 2023. A prompt pattern catalog to enhance prompt engineering with chatgpt. <i>ArXiv</i> , abs/2302.11382.			
683				
684				
685				
686				
687	Ai Ming Yang, Bin Xiao, Bingning Wang, Borong Zhang, Ce Bian, Chao Yin, Chenxu Lv, Da Pan, Dian Wang, Dong Yan, Fan Yang, Fei Deng, Feng Wang, Feng Liu, Guangwei Ai, Guosheng Dong, Hai Zhao, Hang Xu, Hao-Lun Sun, Hongda Zhang, Hui Liu, Jiaming Ji, Jian Xie, Juntao Dai, Kuncheng Fang, Lei Su, Liang Song, Lifeng Liu, Liyun Ru, Luyao Ma, Mang Wang, Mickel Liu, MingAn Lin, Nuolan Nie, Pei Guo, Ruiyang Sun, Zhang Tao, Tianpeng Li, Tianyu Li, Wei Cheng, Weipeng Chen, Xiangrong Zeng, Xiaochuan Wang, Xiaoxi Chen, Xin Men, Xin Yu, Xuehai Pan, Yan-Bin Shen, Yiding Wang, Yiyu Li, Youxin Jiang, Yuchen Gao, Yupeng Zhang, Zenan Zhou, and Zhiying Wu. 2023a. Baichuan 2: Open large-scale language models. <i>ArXiv</i> , abs/2309.10305.			
688				
689				
690				
691				
692				
693				
694				
695				
696				
697				
698				
699				
700				
701				
702	Chengrun Yang, Xuezhi Wang, Yifeng Lu, Hanxiao Liu, Quoc V. Le, Denny Zhou, and Xinyun Chen. 2023b. Large language models as optimizers. <i>ArXiv</i> , abs/2309.03409.			
703				
704				
705				
706	Chengrun Yang, Xuezhi Wang, Yifeng Lu, Hanxiao Liu, Quoc V Le, Denny Zhou, and Xinyun Chen. 2024. Large language models as optimizers. In <i>The Twelfth International Conference on Learning Representations</i> .			
707				
708				
709				
710				