

Rethinking Prompt Optimizers: From Prompt Merits to Optimization

Anonymous ACL submission

Abstract

Prompt optimization (PO) provides a practical way to improve response quality when users lack the time or expertise to manually craft effective prompts. Existing methods typically rely on advanced, large-scale LLMs like GPT-4 to generate optimized prompts. However, due to limited downward compatibility, verbose, and instruction-heavy prompts from advanced LLMs can overwhelm lightweight inference models and degrade response quality. In this work, we rethink prompt optimization through the lens of explicit and interpretable design. We first identify a set of model-agnostic prompt quality merits and empirically validate their effectiveness in enhancing prompt and response quality. We then introduce a merit-guided prompt optimizer (MePO), which is locally deployable and trained on our preference dataset built from merit-guided prompts generated by a lightweight LLM. Unlike prior work, MePO avoids online optimization reliance, reduces cost and privacy concerns, and—by learning clear, interpretable merits—generalizes effectively to both large-scale and lightweight inference models. Experiments demonstrate that MePO achieves better results across diverse tasks and model types, offering a scalable and robust solution for real-world deployment.

1 Introduction

Large language models (LLMs) have achieved impressive results across many NLP tasks (Achiam et al., 2023; Touvron et al., 2023), but their performance remains highly sensitive to prompt phrasing. Although careful prompt crafting can improve output quality, it is often impractical in real-world settings, where users typically lack the time or expertise to manually refine prompts. This has sparked growing interest in automatic prompt optimization (APO) (Liu et al., 2023).

A dominant APO paradigm is discrete automatic prompt optimization, which employs an advanced,

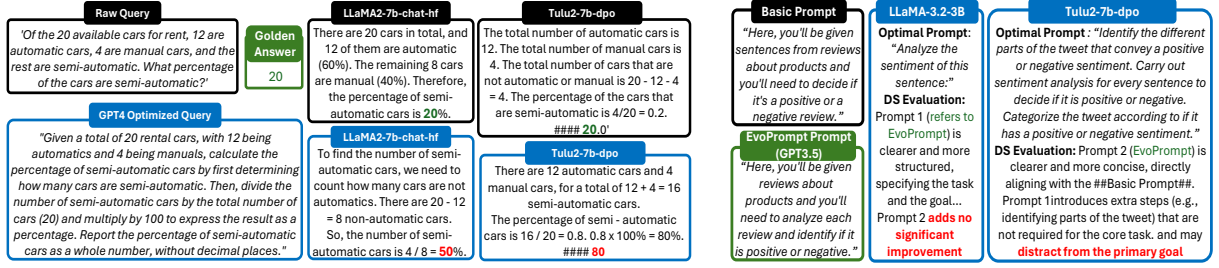
large-scale LLM¹ as the **prompt optimization model** to optimize prompts. These prompts are then used as inputs to either similarly sized (Xiang et al., 2025; Ye et al., 2024) or smaller (Guo et al., 2024) **inference models** for response generation.

One line of work uses API-based models such as GPT-4 for online prompt optimization, typically inserting a task-specific *meta-prompt* before or after each query to reduce cost, while leaving the query itself unoptimized (Xiang et al., 2025; Guo et al., 2024; Ye et al., 2024). However, this approach requires prior task knowledge, limiting its applicability in open-ended, real-world scenarios. Another line explores locally deployable optimizers that optimize *each prompt* and are trained independently of specific tasks, though their training data still originates from prompts generated by advanced online LLMs, implicitly treating these outputs as optimal (Lu et al., 2025; Cheng et al., 2024).

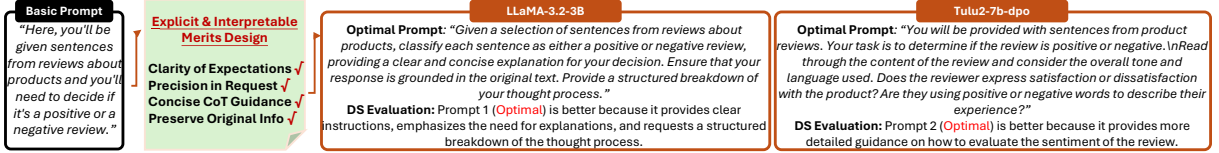
Despite the differences in deployment strategies, both online and local optimization approaches rely on prompts produced by advanced online LLMs as the gold-standard optimal prompts. This design implicitly assumes **downward compatibility**: the notion that prompts generated by large-scale LLMs can be effectively interpreted by lightweight inference models. However, our empirical analysis challenges this assumption. As shown in Fig. 1(a), prompts crafted by advanced, large-scale LLMs often contain **verbose, instruction-heavy** content, which may **overwhelm lightweight models**, causing them to miss the key point or engage in excessive reasoning, ultimately leading to incorrect or irrelevant responses.

The finding on the limited downward compatibility of large-scale LLMs raises a crucial question: *Can lightweight LLMs themselves serve as effective optimization models?*

¹In this work, we categorize LLMs ranging from GPT-4-size to 70B as large-scale, while those with fewer than 13B are denoted as lightweight, following ChunLiu et al. (2024).



(a) Prompts optimized by advanced models can trigger incorrect responses from lightweight inference models. Irrelevant response examples are provided in Appx. A.1. (b) Lightweight LLMs fail under complex optimization algorithms, yielding results equal to or worse than the Basic Prompt.



(c) Lightweight LLMs perform well as PO models with merit guidance.

Figure 1: Empirical analysis of prompt optimization behavior across model scales and optimization algorithms. (a) Raw and GPT-4 optimized prompts are drawn from the GSM8K (Cobbe et al., 2021) and Lu et al. (2025). (b) Basic and EvoPrompt examples, along with the optimization algorithm used, are adapted from Guo et al. (2024). (c) The optimized prompt is generated using our merit-guided instruction (Fig. 14). We present DeepSeek-R1 (DS)’s prompt and response evaluations (Results are consistent with GPT-4o). Further details are provided in Appx.A.

Lightweight models are typically not used for prompt optimization due to their perceived limitations in handling complex optimization processes (Wang et al., 2024; Lu et al., 2025), as further illustrated in Fig. 1(b). However, our research reveals a promising alternative. We find that these models can indeed serve as effective prompt optimizers when guided by explicit and interpretable prompt design merits (Fig. 1(c)). This suggests that, with **clear, learnable structures**, even lightweight LLMs can function as capable optimizers, offering a cost-efficient and locally deployable alternative to online methods, particularly in low-resource or privacy-sensitive settings.

To this end, we propose **MePO**, a **Merit-Guided Prompt Optimization** model. We begin by conducting a systematic empirical analysis to identify four sets of interpretable prompt merits that consistently characterize high-quality prompts, offering actionable insights into effective prompt construction. Guided by these merits, we build a prompt optimization preference dataset using optimized prompts generated by a lightweight LLM—eschewing reliance on advanced online LLMs as in prior work—and use it to train an end-to-end prompt optimization model. At the prompt level, our approach **produces precise and clear merit guidance that can be generated by lightweight prompt optimization models**. At the

response level, our optimized prompts not only enhance the performance of similarly scaled inference models, but also demonstrate **strong downward and upward compatibility**, as the learned prompt merits generalize effectively across both large-scale and lightweight inference models. Besides, constructing the dataset using lightweight models greatly reduces API cost and external dependencies, enabling scalable, privacy-preserving, and locally deployable prompt optimization.

Our main contributions are as follows:

1. We identify and formalize a set of interpretable prompt merits that contribute to both high-quality prompts and responses.
2. We construct a 40k-scale, merit-guided, prompt optimization preference dataset using lightweight LLMs while maintaining high optimization quality.
3. We propose MePO, a lightweight, locally deployable prompt optimization model trained on our merit-guided preference dataset.
4. Evaluations of MePO at both the prompt and response levels show it matches or exceeds existing discrete APO methods and remains effective across inference models of varying capacities.

2 Empirical Analysis: What Merits a Good Prompt?

Prior APO frameworks commonly rely on advanced online LLMs (Guo et al., 2024; Yuksek-

gonul et al., 2024; Cheng et al., 2024) to generate optimized prompts, directly leveraging the high-quality prompt generation capabilities of large-scale models (Pan et al., 2023). However, some studies raise concerns about the transferability of such prompt optimization capabilities to smaller models. Wang et al. (2024) caution that PromptAgent is designed to optimize prompts for state-of-the-art LLMs, yet these expert-level prompts often fail to transfer to smaller models such as GPT-2 or LLaMA2-7b, resulting in substantial performance degradation. Similarly, Lu et al. (2025) observe that small optimizers (13b) fail in difficult prompt optimization tasks.

We argue that these limitations are not solely due to the generation capacity of lightweight LLMs, but rather to the overly **vague or underspecified optimization instructions** used in prior work. Specifically, Wang et al. (2024) adopt the prompt “Given error feedback, generate a better prompt,” and Lu et al. (2025) use “You are an expert in prompt optimization.” Both are overly general and lack concrete guidance. In contrast, we find that when lightweight LLMs are explicitly instructed with detailed prompt patterns, they can generate effective optimized prompts despite their limited capacity (Fig. 1(c)). This observation motivates a key question: *What merits a good prompt?*

To enable lightweight LLMs as effective prompt optimizers, in this section, we analyze response-level and prompt-level characteristics to identify core merits that contribute to high-quality prompts.

2.1 Merit Discovery

Prior studies explore various merits that contribute to prompt effectiveness. Wei et al. (2022) propose that including Chain-of-Thought (CoT) reasoning in prompts enhances performance, while Lampinen et al. (2022) find that Detailed Explanations improve prompt quality. Bsharat et al. (2023) introduce 26 prompting principles, categorized into dimensions such as Prompt Structure and Clarity, and Specificity and Information. Ye et al. (2024) propose three key components for complex reasoning prompts: Detailed Descriptions, Context Specification, and Step-by-Step Reasoning.

Although each of these works contributes valuable insights, they often emphasize different dimensions, resulting in a lack of consensus on what constitutes a high-quality prompt. To further explore this space, we take an empirical approach to uncover the merits of effective prompts through

comparative evaluation.

Empirical Analysis: We randomly select 5000 raw questions from the Alpaca dataset (Taori et al., 2023) and ask a lightweight LLM² to rewrite each question five times. We then use the same LLM as the inference model to generate responses for both the raw and rewritten questions. The quality of each response is scored using DeepSeek-R1 (Guo et al., 2025). Interestingly, we observe that even small changes in a prompt—sometimes just a few words—can lead to significant differences in response scores. This suggests that certain rewrites yield better prompts, which in turn lead to more effective answers.

Based on this observation, we identify raw-rewritten prompt pairs and categorize them as: (1) prompts that led to higher-scoring responses, and (2) those that resulted in lower-scoring responses, with a score difference greater than 4 points. We then ask DeepSeek-R1 to identify the merits that make the higher-scoring prompts more effective compared to the lower-scoring ones.




Figure 2: Key merits of high-performing prompts extracted from DeepSeek-R1 evaluations.

Fig. 2 shows the top 10 merits most frequently associated with high-scoring prompts. Among them, several merits exhibit notably higher frequencies, revealing that high-quality prompts consistently reflect a set of core merits that can be effectively generated by lightweight LLMs. Motivated by this finding and prior insights, we propose the following design merits to guide the construction of optimal prompts:

- **Clarity:** The optimal prompt should set clear, unambiguous expectations for the responder to enable a thorough and accurate reply.
- **Precision:** Use more precise and purposeful language, especially when referring to selecting words or concepts without a fixed pattern.
- **Concise Chain-of-Thought:** Include brief yet contextually rich reasoning or structural cues to

²<https://huggingface.co/Qwen/Qwen2.5-7B-Instruct>

guide the responder’s thought process, while remaining focused and concise.

- **Preserve Original Information:** Focus on the original prompt, ensuring no information or intent is lost or omitted in the transformation.

Guided by the discovered merits, we optimize the prompts into their merit-guided versions using a lightweight LLM under the instruction in Fig. 14. **Response-Level Evaluation:** To assess whether merit-guided prompts improve response quality, we compare responses generated from merit-guided prompts against those from raw prompts across two datasets (2,000 samples each), with scores assigned by DeepSeek-R1.

- **Alpaca Dataset (30k)**³: Raw responses are generated by text-davinci-003

- **BPO Dataset (13.9k)**⁴: Raw responses are sourced from human-preferred answers.




Figure 3: Response-level win rate comparison. ‘Win’ indicates that the merit-guided prompt’s response received a higher score than that from the raw prompt.

As shown in Fig. 3, prompts optimized using our designed merits—despite being applied to a lightweight model—can match or even outperform GPT-3.5-generated responses (Alpaca) and significantly beat human-preferred responses (BPO).

Prompt-Level Evaluation: To move beyond response quality and directly assess prompt quality—a dimension largely overlooked in prior work—we randomly interleaved optimized and raw prompts and asked DeepSeek-R1 to evaluate them in isolation (i.e., without seeing the generated responses). Across 2000 prompt comparisons per dataset, optimized prompts were preferred in 95.75% of cases on BPO and 97.9% on Alpaca. These results reinforce the effectiveness of our proposed prompt merits and demonstrate that prompt quality can be reliably evaluated independently of generated outputs.

Summary: Our evaluations demonstrate that optimized prompts grounded in our proposed interpretable merits not only yield stronger responses but also exhibit higher prompt quality. More-

³<https://huggingface.co/datasets/tatsu-lab/alpaca>

⁴<https://huggingface.co/datasets/THUDM/BPO>

over, the results show that effective prompt optimization is not exclusive to advanced, large-scale LLMs such as GPT-4. With merit-based guidance, lightweight LLMs can also serve as capable prompt optimizers, offering scalable and privacy-friendly alternatives for real-world deployment.

The detailed implementation of merit discovery process is provided in Appx.B.

3 Method: From Merits to Optimizer

To incorporate the proposed merits into the APO pipeline, we apply them at the optimization stage as supervision signals for preference learning⁵.

Our objective is to end-to-end optimize arbitrary input prompts P_{silver} into refined prompts P_{golden} , enabling inference models to generate higher-quality responses. To achieve this, we introduce **MePO**, a lightweight LLM-based prompt optimization model designed for local deployment. MePO is trained under a preference learning framework, supervised by our proposed dataset, where prompts are constructed according to the merits identified in Sec.2. An overview of MePO is shown in Fig.4.

3.1 Constructing the Merit-Aligned Prompt Preference Dataset

To train MePO, we construct an API-free Prompt Optimization Preference (POP) dataset grounded in our discovered merits. Following Cheng et al. (2024); Lu et al. (2025), we adopt Alpaca and BPO as our base datasets. The instruction field in Alpaca and the prompt field in BPO are treated as raw prompts P_{silver} , which are then refined by a lightweight LLM M_{ref} ² using the defined merits to generate P_{golden} , eliminating dependence on online LLM optimization.

We then use M_{ref} to generate responses \hat{R}_{golden} from each optimized prompt P_{golden} , and retrieve the corresponding raw responses \hat{R}_{silver} from Alpaca’s output and BPO’s bad_res. Each response is then evaluated by DeepSeek-R1, and we retain only those pairs that satisfy two criteria: (1) \hat{R}_{golden} receives a higher evaluation score than \hat{R}_{silver} , and (2) the score of \hat{R}_{golden} exceeds 8⁶. The retained

⁵Alternatively, the merits could be used in conjunction with response feedback to iteratively guide the optimization process through the inference model. In this work, we focus on the optimization stage and leave integration with inference feedback for future exploration.

⁶According to our scoring rubric in Fig. 13, responses with scores above 8 are considered accurate and well-aligned with the prompt.

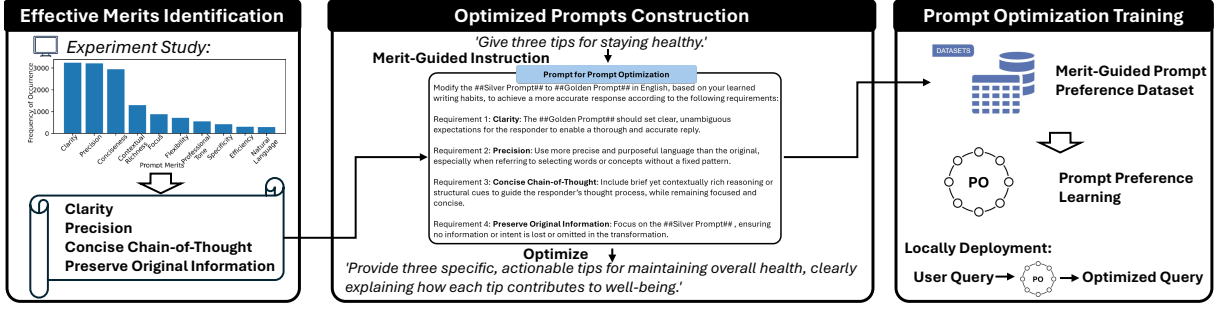


Figure 4: Overall stages of merit-guided prompt optimization (MePO).

samples are relabeled as R_{golden} and R_{silver} , forming a structured 4-tuple used in training:

$$(P_{\text{silver}}, P_{\text{golden}}, R_{\text{silver}}, R_{\text{golden}}).$$

Since Alpaca and BPO contain relatively well-formed prompts—either produced by text-davinci-003 or curated by humans—we further simulate real-world usage by randomly selecting 10% of P_{silver} and intentionally degrading them using a base model. These lower-quality prompts better reflect common user inputs and are paired with their corresponding P_{golden} to enhance data diversity.

Prompt Source Statistics		Alpaca		BPO	
Category		Naive	Degraded	Naive	Degraded
# 4-Tuples		25,526	3,000	10,225	1,400
Total samples: 40,151					
Validation (Win Rate)		GPT-4o (5k)		DeepSeek-R1 (40k)	
Response Win Rate		94.82%		100.00%	
Prompt Win Rate		97.68%		97.62%	

Table 1: Overview of our POP dataset. The top section reports the number of prompt–response 4-tuples derived from Alpaca and BPO. The bottom section shows validation results from DeepSeek-R1 and GPT-4o, with *win rate* denoting the proportion where R_{golden} or P_{golden} outperforms its silver counterpart.

After filtering and aggregation, the final POP dataset comprises 40k samples. Table 1 summarizes the dataset composition and presents quality validation scores from both DeepSeek-R1 and GPT-4o, showing consistently strong preference for the optimized prompts across both response- and prompt-level evaluations.

3.2 Training the Merit-Guided Prompt Optimizer

To guide the model toward generating more preferred prompts, we train the optimizer using the Direct Preference Optimization (DPO) framework (Rafailov et al., 2023). The objective is to

learn a prompt optimizer that prefers merit-guided prompts yielding higher-quality responses.

In our setup, each training instance includes a *chosen* prompt P_{golden} that yields a higher-quality response and a *rejected* prompt P_{silver} with a lower-quality response.

The DPO objective is formulated as:

$$\mathcal{L}_{\text{DPO}}(M_o; M_{\text{ref}}) = -\mathbb{E}_{(x, c, r) \sim \mathcal{D}} [\log \sigma(\beta \cdot \Delta)]$$

$$\Delta = \log \frac{M_o(c | x)}{M_{\text{ref}}(c | x)} - \log \frac{M_o(r | x)}{M_{\text{ref}}(r | x)}$$

where M_{ref} serves as the reference model and M_o is the optimization model (M_{ref} with adapters).

Further implementation details on dataset construction and optimizer training are in Appx.C.

4 Experiments

We evaluate the effectiveness of MePO from the following perspectives: (1) Comparison with SOTA local prompt optimizers (Sec. 4.2); (2) Ablation study with training-free prompting (Sec. 4.3); (3) Upward compatibility in larger inference models (Sec. 4.4); (4) Comparison with online-based meta-prompts (Sec. 4.5) and (5) Effectiveness analysis of optimized prompts (Sec. 4.6).

4.1 Setups

MePO is trained using Qwen2.5-7B-Instruct², selected for its strong generation ability among similarly sized models (Qwen et al., 2025). This choice ensures that the optimizer is built on a lightweight LLM that maximizes performance within its size class—allowing efficient deployment without sacrificing generation quality.

Datasets We evaluate our method on five benchmark datasets: ARC-Easy, ARC-Challenge (Clark et al., 2018), GSM8K (Cobbe et al., 2021), BigBench-Hard (BBH) (Suzgun et al., 2023), and PiQA (Bisk et al., 2020). Following Lu et al.

Inference Model	Optimizer	ARC-Easy	ARC-Challenge	GSM8K	BBH	PiQA	Avg.
Qwen2-7b	-	80.68	65.19	76.88	49.34	80.20	70.46
	Inference Model	81.22	66.21	79.01	52.68	81.34	72.09
	BPO	80.89	66.38	77.38	53.07	82.64	72.07
	FIPO	82.37	67.49	82.71	52.74	82.48	73.56
	MePO	83.33	68.52	83.12	54.35	83.46	74.56
Tulu2-7b-dpo	-	45.37	26.54	30.63	35.67	70.02	41.65
	Inference Model	46.40	29.35	32.93	38.27	70.78	43.55
	BPO	47.98	31.57	32.08	39.33	71.60	44.51
	FIPO	50.55	36.95	33.66	39.56	72.63	46.67
	MePO	55.05	38.14	35.18	43.25	73.60	49.04
LLaMA2-7b-chat-hf	-	35.40	29.27	15.06	34.17	49.95	32.77
	Inference Model	36.74	29.52	16.27	36.65	51.41	34.12
	BPO	38.30	30.89	19.85	39.60	52.56	36.24
	FIPO	36.24	29.44	17.03	39.14	51.58	34.69
	MePO	39.86	31.74	25.25	41.97	55.33	38.83
LLaMA3-8b-instruct	-	46.72	41.31	55.04	38.76	65.56	49.48
	Inference Model	52.02	42.32	57.87	40.97	67.19	52.07
	BPO	50.88	44.37	58.39	42.68	71.81	53.63
	FIPO	49.12	44.97	56.81	42.66	75.56	53.82
	MePO	55.01	51.88	59.59	45.52	76.66	57.73

Table 2: Performance comparison of prompt optimization methods across multiple datasets and inference models. ‘-’ denotes the original unoptimized prompt. ‘Inference Model’ refers to using the inference model itself as the optimizer (based on the instruction format in Fig. 11). Detailed BBH task-level results are provided in Appx.D.3.

(2025), we adopt 3-shot for GSM8K, BBH, and PiQA, and zero-shot for ARC, with instructions for answer extraction. Prompt optimization is applied only to test query, while others remain unchanged. **Baselines** We compare MePO with two state-of-the-art local prompt optimization models: BPO (Cheng et al., 2024) and FIPO (Lu et al., 2025). We evaluate optimizers’ performance across three inference model families: Qwen (Yang et al., 2024), Tulu (Iverson et al., 2023) and LLaMA (Touvron et al., 2023). We also use the inference model itself as an optimization baseline, motivated by the inherent alignment between a model and its own preferences, as leveraging an LLM’s natural capability to refine prompts for itself better reflects its aligned outputs (Xiang et al., 2025).

Additional implementation details and dataset descriptions are provided in Appx.D.

4.2 Main Results

We present the performance of each optimizer, including our proposed MePO, on different inference models across five datasets. The results, summarized in Table 2, lead to the following conclusions: **Efficacy** MePO consistently outperforms all baselines across datasets and task models, with average accuracy gains of 1% on Qwen2-7B, 2.37% on Tulu2-7B, 2.59% on LLaMA2-7B, and 3.91% on LLaMA3-8B. Notable improvements are observed on ARC-Challenge with LLaMA3-8B (+6.91%),

GSM8K with LLaMA2-7B (+5.4%), and ARC-Easy and BBH with Tulu2-7B (+4.5%, +3.69%).

Robustness Inference models often perform better when the optimizer shares the same architecture, reflecting alignment with model-specific preferences. For example, in the BPO–FIPO comparison, BPO (based on LLaMA2) performs better when responding with LLaMA2-7B, while FIPO (based on Tulu2) excels with Tulu2-7B. However, both optimizers show misalignment on models like LLaMA3 and Qwen2. On ARC-Easy, inference model optimization outperforms BPO (+1.14% with LLaMA3, +0.33% with Qwen2) and FIPO (+2.90% with LLaMA3). Similarly, on GSM8K, the inference model surpasses BPO (+1.63% with Qwen2) and FIPO (+1.06% with LLaMA3). In contrast, MePO consistently outperforms inference model optimization across all datasets, indicating that its learned prompt merits are broadly interpretable across architectures, demonstrating both robustness and strong generalization capability.

Downward Compatibility FIPO shows performance degradation compared to the inference model optimizer on LLaMA2-7b (-0.5% on ARC-Easy, -0.08% on ARC-Challenge) and LLaMA3-8b (-1.06% on GSM8K, -2.9% on ARC-Easy). As FIPO is a 70b optimizer trained on prompts optimized by GPT-4, its outputs may be too complex for smaller inference models to interpret—indicating a lack of downward compatibility.

Task Model	Optimizer	ARC-Easy	ARC-Challenge	GSM8K	BBH	PiQA	Avg.
Qwen2-7b	Training Free	82.79	68.17	81.52	52.93	82.32	73.55
	PattoPO	83.33	68.52	83.12	54.35	83.46	74.56
Tulu2-7b-dpo	Training Free	49.58	34.30	32.93	40.96	72.60	46.22
	PattoPO	55.05	38.14	35.18	43.25	73.60	49.04
LLaMA2-7b-chat-hf	Training Free	36.91	30.55	20.91	39.88	53.10	36.27
	PattoPO	39.86	31.74	25.25	41.97	55.33	38.83
LLaMA3-8b-instruct	Training Free	52.36	50.77	58.25	42.80	71.69	55.17
	PattoPO	55.01	51.88	59.59	45.52	76.66	57.73

Table 3: Comparison between training-free merit-guided prompts and MePO. Detailed BBH results are in Appx.D.3.

In contrast, MePO, while built upon a relatively strong base model compared to LLaMA2-7b and LLaMA3-8b (Qwen et al., 2025), maintains strong downward compatibility, consistently achieves better performance. More analyses on smaller inference models are in Appx.D.1.

4.3 Ablation Study: Training-Free Comparison

As shown in Table 1, we retain 25.5k Alpaca and 10k BPO samples by discarding cases where merit-guided prompts fail to outperform their raw counterparts. This ensures MePO is trained only on reliably successful examples, promoting robust learning of merit application.

To isolate the effect of learning, we compare MePO with a training-free baseline where the inference model itself rewrites raw prompts using static merit-guided templates (Fig. 14), and then generates responses from the rewritten prompts. This comparison tests whether learning to apply merits is more effective than applying them heuristically.

As shown in Table 3, MePO consistently outperforms the training-free approach across all datasets and model backbones. Notably, average accuracy improves by 2.38% on Tulu2 and over 2.59% on LLaMA2 and 3, demonstrating the value of end-to-end optimization over training-free prompting.

4.4 Case Study: Upward Compatibility

To evaluate MePO’s upward compatibility, we examine whether its optimized prompts improve response quality in large-scale inference models.

As shown in Table 4, applying MePO prompts to larger inference LLMs consistently matches or improves performance across all datasets. These results demonstrate that MePO, despite being lightweight, produces clear and concise reasoning guidance that large-scale LLMs can effectively interpret, highlighting its strong upward compatibility and potential for general applicability.

	LLaMA2-13b -chat-hf		Llama-3.3-70B -Instruct		Tulu2-70b -dpo	
	-	MePO	-	MePO	-	MePO
ARC-Easy	49.83	55.93	95.03	92.34	80.47	85.23
ARC-Challenge	45.05	52.73	92.24	91.04	58.87	60.49
GSM8K	28.73	35.97	91.74	92.49	61.41	64.67
BBH	39.49	46.74	66.25	68.25	51.63	53.07
PiQA	56.13	60.99	87.87	89.93	77.75	80.14

Table 4: Response evaluation of MePO prompts in larger LLMs. Detailed BBH results are in Table 9 and 14.

4.5 Case Study: Meta-Prompt Comparison

To further validate MePO’s effectiveness, we compare it on GSM8K with three online-based methods: APO (Pryzant et al., 2023), APE (Zhou et al., 2022), and PE2 (Ye et al., 2024). These methods enhance reasoning by appending an optimized meta-prompt after the query, keeping the original query unchanged while MePO directly optimizes each query. As shown in Table 5, MePO outperforms all baselines, indicating that per-query merit-driven optimization is more effective than applying a fixed meta-prompt across diverse inputs.

Optimizer	Tulu2-7b-dpo	LLaMA2-7b-chat-hf
APO	29.19	22.97
Iterative APE	29.72	20.92
PE2	32.15	22.37
MePO	35.18	25.25

Table 5: Comparison between meta-prompts and MePO. Details of the applied prompts are in Appx.D.2.

4.6 Case Study: Interpretability of MePO

To understand the effectiveness of MePO’s optimized prompts, we conduct a qualitative analysis on samples from the GSM8K dataset.

Example 1

Raw: Carla is downloading a 200 GB file. Normally she can download 2 GB/minute, but 40% of the way through the download, Windows forces a restart to install updates, which takes 20 minutes. Then Carla has to restart the download from the

beginning. How long does it take to download the file?

MePO: Carla is attempting to download a 200 GB file at a rate of 2 GB per minute. At 40% completion, her download is interrupted for 20 minutes to allow for a Windows update installation. After this interruption, she must restart the entire download from the beginning. Calculate the total time required to complete the file download, including the initial download time, the update interruption, and the subsequent restart.

Observation: The optimized prompt offers clearer structure and explicitly outlines the three components required for correct reasoning. This clarity facilitates more accurate downstream responses.

Example 2

Raw: Shiela bought five cell phones for \$150 each for a 3-month installment. A 2% interest will be charged for each unit. How much will Shiela pay each month for 3 months?

MePO: Shiela purchased five cell phones at \$150 each, with a 3-month installment plan that includes a 2% total interest rate on each phone. Calculate the total monthly payment Shiela will make over the 3-month period, considering the interest applied to each phone separately.

Observation: The raw prompt is ambiguous—it’s unclear whether the 2% interest is monthly or total. The MePO version resolves this ambiguity by clearly specifying a 2% total interest rate, enabling the model to reason correctly.

These examples illustrate how MePO’s optimized prompts reflect our defined merits—clarity, precision, and concise CoT guidance—to reduce ambiguity and improve task performance.

5 Related Work

Automatic Prompt Optimization (APO) provides a practical alternative to fine-tuning by optimizing prompts (Yang et al.; Zhu et al., 2023). Beyond exploration-based methods search over predefined prompt pools (Ma et al., 2023; Shi et al.), recent works favor discrete exploitation-based methods that directly generate or revise prompts without relying on the quality of initial pools, making them more flexible and better suited for generative tasks beyond fixed-answer classification.

Online-Based Discrete APO Several works perform discrete prompt optimization using high-capacity, API-based LLMs as optimization models. For example, Zhou et al. (2022) use InstructGPT;

Wang et al. (2024); Guo et al. (2024); Pryzant et al. (2023) rely on GPT-3.5; Xiang et al. (2025) adopt Claude 3.5; and Yuksekgonul et al. (2024) utilize GPT-4o. While effective, these approaches depend on proprietary APIs, introducing practical concerns around cost, inference latency, and privacy risks.

Local Discrete APO Locally trainable prompt optimizers have been explored for general use. FIPO (Lu et al., 2025) trains a Tulu2-based model using a GPT-3.5/4-generated preference dataset. BPO (Cheng et al., 2024) uses a LLaMA-based optimizer trained on ChatGPT-optimal prompts. MAPO (Chen et al., 2023) adopts a model-adaptive strategy with GPT-3.5-generated prompts but requires a separate warm-up set per task, limiting generalization. In contrast, our model uses merit-guided prompts from a lightweight LLM, achieving strong performance without API reliance.

Prompt Merits Exploration To improve prompt quality, Arora et al. show that prompts in open-ended QA formats outperform restrictive ones. Wei et al. (2022) introduce intermediate reasoning steps in prompts. Zhou et al. (2023) propose decomposing prompts into simpler subcomponents. Besides, several studies provide design principles for prompt optimization (Bsharat et al., 2023; Ye et al., 2024). However, these merits are largely proposed heuristically and validated through downstream experiments, leading to inconsistencies in perspective. In contrast, we derive prompt merits by prompt–response empirical analysis, grounding them in measurable improvements.

6 Conclusion

In this work, we introduce MePO, a lightweight, locally deployable prompt optimization model trained under a merit-guided preference framework. Empirical analyses are conducted to discover prompt merits—clarity, precision, and concise chain-of-thought—that contribute to high-quality prompts, resulting in a prompt preference dataset proposed using the lightweight LLMs. Experimental results show that MePO not only generates structurally clear and precise prompts, but also exhibits strong downward and upward compatibility, maintaining robust performance across variously scaled inference models. Our findings demonstrate that, with well-defined optimization merits, even lightweight LLMs can serve as effective prompt optimizers—enabling scalable, cost-efficient, and privacy-friendly deployment in real-world settings.

7 Limitations

While MePO demonstrates strong performance across diverse tasks and model scales, several limitations remain: (1) Lack of Interactive Feedback: MePO currently operates as a one-shot prompt optimizer, without incorporating iterative signals from users or inference model feedback. In real-world deployments, both user interactions and model responses can provide valuable guidance for continual prompt refinement. Integrating MePO into an interactive, feedback-driven optimization loop remains a promising direction for future work. (2) Limited Model Adaptation: Although MePO exhibits robustness across different inference models, further gains may be possible by aligning the base architectures of the optimization and inference models. While our results show that MePO performs well even under architectural mismatch, explicitly training model-adaptive optimizers could improve performance by leveraging shared internal representations.

References

Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.

Simran Arora, Avani Narayan, Mayee F Chen, Laurel Orr, Neel Guha, Kush Bhatia, Ines Chami, and Christopher Re. Ask me anything: A simple strategy for prompting language models. In *The Eleventh International Conference on Learning Representations*.

Yonatan Bisk, Rowan Zellers, Jianfeng Gao, Yejin Choi, et al. 2020. Piqa: Reasoning about physical commonsense in natural language. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pages 7432–7439.

Sondos Mahmoud Bsharat, Aidar Myrzakhan, and Zhiqiang Shen. 2023. Principled instructions are all you need for questioning llama-1/2, gpt-3.5/4. *arXiv preprint arXiv:2312.16171*, 3.

Yuyan Chen, Zhihao Wen, Ge Fan, Zhengyu Chen, Wei Wu, Dayiheng Liu, Zhixu Li, Bang Liu, and Yanghua Xiao. 2023. Mapo: Boosting large language model performance with model-adaptive prompt optimization. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 3279–3304.

Jiale Cheng, Xiao Liu, Kehan Zheng, Pei Ke, Hongning Wang, Yuxiao Dong, Jie Tang, and Minlie Huang. 2024. Black-box prompt optimization: Aligning large language models without model training. In

Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), pages 3201–3219.

ChunLiu ChunLiu, Hongguang Zhang, Kainan Zhao, Xinghai Ju, and Lin Yang. 2024. Llmembd: Rethinking lightweight llm’s genuine function in text classification. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 7994–8004.

Peter Clark, Isaac Cowhey, Oren Etzioni, Tushar Khot, Ashish Sabharwal, Carissa Schoenick, and Oyvind Tafjord. 2018. Think you have solved question answering? try arc, the ai2 reasoning challenge. *arXiv:1803.05457v1*.

Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, et al. 2021. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*.

Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.

Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, et al. 2025. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*.

Qingyan Guo, Rui Wang, Junliang Guo, Bei Li, Kaitao Song, Xu Tan, Guoqing Liu, Jiang Bian, and Yujiu Yang. 2024. Connecting large language models with evolutionary algorithms yields powerful prompt optimizers. In *The Twelfth International Conference on Learning Representations*.

Hamish Ivison, Yizhong Wang, Valentina Pyatkin, Nathan Lambert, Matthew Peters, Pradeep Dasigi, Joel Jang, David Wadden, Noah A. Smith, Iz Beltagy, and Hannaneh Hajishirzi. 2023. Camels in a changing climate: Enhancing lm adaptation with tulu 2. *Preprint*, arXiv:2311.10702.

Andrew Lampinen, Ishita Dasgupta, Stephanie Chan, Kory Mathewson, Mh Tessler, Antonia Creswell, James McClelland, Jane Wang, and Felix Hill. 2022. Can language models learn from explanations in context? In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 537–563.

Pengfei Liu, Weizhe Yuan, Jinlan Fu, Zhengbao Jiang, Hiroaki Hayashi, and Graham Neubig. 2023. Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing. *ACM computing surveys*, 55(9):1–35.

Junru Lu, Siyu An, Min Zhang, Yulan He, Di Yin, and Xing Sun. 2025. Fipo: Free-form instruction-oriented prompt optimization with preference dataset and modular fine-tuning schema. In *Proceedings of*

698	<i>the 31st International Conference on Computational Linguistics</i> , pages 11029–11047.	
699		
700	Huan Ma, Changqing Zhang, Yatao Bian, Lemao Liu,	
701	Zhirui Zhang, Peilin Zhao, Shu Zhang, Huazhu Fu,	
702	Qinghua Hu, and Bingzhe Wu. 2023. Fairness-	
703	guided few-shot prompting for large language mod-	
704	els. <i>Advances in Neural Information Processing Sys-</i>	
705	<i>tems</i> , 36:43136–43155.	
706	Alexander Pan, Jun Shern Chan, Andy Zou, Nathaniel	
707	Li, Steven Basart, Thomas Woodside, Hanlin Zhang,	
708	Scott Emmons, and Dan Hendrycks. 2023. Do the	
709	rewards justify the means? measuring trade-offs be-	
710	tween rewards and ethical behavior in the machiavelli	
711	benchmark. In <i>International conference on machine</i>	
712	<i>learning</i> , pages 26837–26867. PMLR.	
713	Reid Pryzant, Dan Iter, Jerry Li, Yin Lee, Chenguang	
714	Zhu, and Michael Zeng. 2023. Automatic prompt op-	
715	timization with “gradient descent” and beam search.	
716	In <i>Proceedings of the 2023 Conference on Empiri-</i>	
717	<i>cal Methods in Natural Language Processing</i> , pages	
718	7957–7968.	
719	Qwen, :, An Yang, Baosong Yang, Beichen Zhang,	
720	Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li,	
721	Dayiheng Liu, Fei Huang, Haoran Wei, Huan Lin,	
722	Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang,	
723	Jiaxi Yang, Jingren Zhou, Junyang Lin, Kai Dang,	
724	Keming Lu, Keqin Bao, Kexin Yang, Le Yu, Mei Li,	
725	Mingfeng Xue, Pei Zhang, Qin Zhu, Rui Men, Runji	
726	Lin, Tianhao Li, Tianyi Tang, Tingyu Xia, Xingzhang	
727	Ren, Xuancheng Ren, Yang Fan, Yang Su, Yichang	
728	Zhang, Yu Wan, Yuqiong Liu, Zeyu Cui, Zhenru	
729	Zhang, and Zihan Qiu. 2025. Qwen2.5 technical	
730	report . <i>Preprint</i> , arXiv:2412.15115.	
731	Rafael Rafailov, Archit Sharma, Eric Mitchell, Christo-	
732	pher D Manning, Stefano Ermon, and Chelsea Finn.	
733	2023. Direct preference optimization: Your lan-	
734	guage model is secretly a reward model. <i>Advances in</i>	
735	<i>Neural Information Processing Systems</i> , 36:53728–	
736	53741.	
737	Chengshuai Shi, Kun Yang, Zihan Chen, Jundong Li,	
738	Jing Yang, and Cong Shen. Efficient prompt opti-	
739	mization through the lens of best arm identification.	
740	In <i>The Thirty-eighth Annual Conference on Neural</i>	
741	<i>Information Processing Systems</i> .	
742	Mirac Suzgun, Nathan Scales, Nathanael Schärli, Se-	
743	bastian Gehrmann, Yi Tay, Hyung Won Chung,	
744	Aakanksha Chowdhery, Quoc Le, Ed Chi, Denny	
745	Zhou, et al. 2023. Challenging big-bench tasks and	
746	whether chain-of-thought can solve them. In <i>Find-</i>	
747	<i>ings of the Association for Computational Linguistics:</i>	
748	<i>ACL 2023</i> , pages 13003–13051.	
749	Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann	
750	Dubois, Xuechen Li, Carlos Guestrin, Percy Liang,	
751	and Tatsunori B. Hashimoto. 2023. Stanford alpaca:	
752	An instruction-following llama model. https://	
753	github.com/tatsu-lab/stanford_alpaca .	
	Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier	754
	Martinet, Marie-Anne Lachaux, Timothée Lacroix,	755
	Baptiste Rozière, Naman Goyal, Eric Hambro,	756
	Faisal Azhar, et al. 2023. Llama: Open and effi-	757
	cient foundation language models. <i>arXiv preprint</i>	758
	<i>arXiv:2302.13971</i> .	759
	Xinyuan Wang, Chenxi Li, Zhen Wang, Fan Bai, Hao-	760
	tian Luo, Jiayou Zhang, Nebojsa Jojic, Eric Xing, and	761
	Zhiting Hu. 2024. Promptagent: Strategic planning	762
	with language models enables expert-level prompt op-	763
	timization . In <i>The Twelfth International Conference</i>	764
	<i>on Learning Representations</i> .	765
	Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten	766
	Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou,	767
	et al. 2022. Chain-of-thought prompting elicits rea-	768
	soning in large language models. <i>Advances in neural</i>	769
	<i>information processing systems</i> , 35:24824–24837.	770
	Jinyu Xiang, Jiayi Zhang, Zhaoyang Yu, Fengwei Teng,	771
	Jinhao Tu, Xinbing Liang, Sirui Hong, Chenglin Wu,	772
	and Yuyu Luo. 2025. Self-supervised prompt opti-	773
	mization. <i>arXiv preprint arXiv:2502.06855</i> .	774
	An Yang, Baosong Yang, Binyuan Hui, Bo Zheng,	775
	Bowen Yu, Chang Zhou, Chengpeng Li, Chengyuan	776
	Li, Dayiheng Liu, Fei Huang, Guanting Dong, Hao-	777
	ran Wei, Huan Lin, Jialong Tang, Jialin Wang,	778
	Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin	779
	Ma, Jianxin Yang, Jin Xu, Jingren Zhou, Jinze Bai,	780
	Jinzheng He, Junyang Lin, Kai Dang, Keming Lu, Ke-	781
	qin Chen, Kexin Yang, Mei Li, Mingfeng Xue, Na Ni,	782
	Pei Zhang, Peng Wang, Ru Peng, Rui Men, Ruize	783
	Gao, Runji Lin, Shijie Wang, Shuai Bai, Sinan Tan,	784
	Tianhang Zhu, Tianhao Li, Tianyu Liu, Wenbin Ge,	785
	Xiaodong Deng, Xiaohuan Zhou, Xingzhang Ren,	786
	Xinyu Zhang, Xipin Wei, Xuancheng Ren, Xuejing	787
	Liu, Yang Fan, Yang Yao, Yichang Zhang, Yu Wan,	788
	Yunfei Chu, Yuqiong Liu, Zeyu Cui, Zhenru Zhang,	789
	Zhifang Guo, and Zhihao Fan. 2024. Qwen2 techni-	790
	cal report . <i>Preprint</i> , arXiv:2407.10671.	791
	Chengrun Yang, Xuezhi Wang, Yifeng Lu, Hanxiao Liu,	792
	Quoc V Le, Denny Zhou, and Xinyun Chen. Large	793
	language models as optimizers. In <i>The Twelfth Inter-</i>	794
	<i>national Conference on Learning Representations</i> .	795
	Zhenyu Yang, Yuhang Hu, Zemin Du, Dizhan Xue,	796
	Shengsheng Qian, Jiahong Wu, Fan Yang, Weim-	797
	ing Dong, and Changsheng Xu. 2025. SVBench: A	798
	benchmark with temporal multi-turn dialogues for	799
	streaming video understanding . In <i>The Thirteenth In-</i>	800
	<i>ternational Conference on Learning Representations</i> .	801
	Qinyuan Ye, Mohamed Ahmed, Reid Pryzant, and	802
	Fereshte Khani. 2024. Prompt engineering a prompt	803
	engineer. In <i>Findings of the Association for Compu-</i>	804
	<i>tational Linguistics ACL 2024</i> , pages 355–385.	805
	Mert Yuksekogonul, Federico Bianchi, Joseph Boen,	806
	Sheng Liu, Zhi Huang, Carlos Guestrin, and James	807
	Zou. 2024. Textgrad: Automatic "differentiation" via	808
	text. <i>arXiv preprint arXiv:2406.07496</i> .	809

Denny Zhou, Nathanael Schärli, Le Hou, Jason Wei, Nathan Scales, Xuezhi Wang, Dale Schuurmans, Claire Cui, Olivier Bousquet, Quoc V Le, and Ed H. Chi. 2023. [Least-to-most prompting enables complex reasoning in large language models](#). In *The Eleventh International Conference on Learning Representations*.

Yongchao Zhou, Andrei Ioan Muresanu, Ziyen Han, Keiran Paster, Silviu Pitis, Harris Chan, and Jimmy Ba. 2022. Large language models are human-level prompt engineers. In *The Eleventh International Conference on Learning Representations*.

Kaijie Zhu, Jindong Wang, Jiaheng Zhou, Zichen Wang, Hao Chen, Yidong Wang, Linyi Yang, Wei Ye, Yue Zhang, Neil Zhenqiang Gong, et al. 2023. Prompt-bench: Towards evaluating the robustness of large language models on adversarial prompts. *arXiv e-prints*, pages arXiv–2306.

A Empirical Study: Prompt Optimization Across Model Scales

A.1 Evaluating GPT-4 Optimized Prompts on Smaller Inference Models

To evaluate the downward compatibility of large-scale optimization models with lightweight inference models, beyond the results shown in Fig. 1(a), we conduct additional experiments that yield irrelevant responses.

We adopt the optimized prompts from Lu et al. (2025), which are derived from a GPT-4 optimization dataset and used to train a Tulu2-70B-DPO model. We then evaluate the effectiveness of these prompts on smaller inference models (LLaMA2-7B and Tulu2-7B-DPO) by generating responses and scoring them with DeepSeek-R1 (see Fig. 13 for scoring setup).

LLaMA2-7B-chat-hf Response:

*"Sure, here's a traditional haiku about the ocean:
Waves roll in, peaceful*

Soothing sounds, endless blue

Lullaby of sea

In this haiku, I wanted to capture the peaceful and soothing quality of ocean waves, much like a lullaby. The image of waves rolling in and out of the shore is a common one in haiku poetry, and I wanted to emphasize the sense of endlessness and tranquility that it evokes. The use of the word 'peaceful' in the first line sets the tone for the rest of the poem, and the image of 'endless blue' in the second line conveys the vastness and depth of the ocean. The final line, 'Lullaby of sea,' adds a sense of melody and rhythm to the poem, which is fitting given the soothing nature of ocean waves."

DeepSeek-R1 Evaluation: The response does not include irrelevant content, but the **additional explanatory text is somewhat unnecessary** for the task.

Tulu2-7B-DPO Response:

"Haiku is a traditional form of Japanese poetry that consists of three lines with a syllable pattern of 5-7-5. It is a concise and evocative form that aims to capture a moment or emotion in a few words. A lullaby, on the other hand, is a song that is typically sung to put children to sleep. It often has a soothing and calming melody, and the lyrics are usually simple and repetitive. Here's my traditional haiku about the ocean:

Gentle ocean waves

Crash upon the shore in peace

Lullaby of the sea

In this haiku, I've captured the peaceful and soothing quality of ocean waves, as they gently crash upon the shore. The repetition of the "ocean waves" phrase in the first and third lines creates a sense of continuity, emphasizing the endless cycle of the ocean's"

DeepSeek-R1 Evaluation: High accuracy in the haiku itself but a minor deduction for the slightly **off-topic** preface.

These examples show that prompts crafted by powerful optimization models may not be directly interpretable by smaller models due to verbosity or overly rich CoT structure, leading to unnecessary or off-topic outputs.

A.2 Evaluating Lightweight LLMs with Complex Optimization Algorithms

We evaluate whether lightweight LLMs can follow optimization instructions designed for advanced, large-scale models. Specifically, we adopt the Differential Evolution (DE) algorithm from Guo et al. (2024), which uses EvoPrompt instructions to guide prompt generation. We provide both the EvoPrompt prompt and the DE-generated optimal prompt as input to lightweight LLMs, and compare their quality using DeepSeek-R1 (see Fig. 5).

As shown in Fig. 1(b), DeepSeek-R1 evaluations report that the EvoPrompt prompt is "clearer and more structured," directly aligning with the task, while the prompt generated by the lightweight LLM "introduces unnecessary steps that may distract from the core goal." These results indicate that GPT-3.5, as used in EvoPrompt, produces higher-quality prompts than those generated by LLaMA-3.2-3B or Tulu2-7b-dpo under the same al-

gorithm—suggesting that lightweight LLMs struggle to serve as effective optimizers when guided by complex optimization instructions.

Evaluating Merits-Guided Prompt Optimization with Lightweight LLMs. We further test whether lightweight LLMs can act as effective optimization models when guided by clear, interpretable merits. Using the Basic Prompt from Guo et al. (2024), we apply our proposed merit-based instruction (see Fig. 14) to generate optimized prompts. We then evaluate the difference between Basic and Optimized Prompts using DeepSeek-R1 with the pairwise setup shown in Fig. 9. The same evaluation is conducted with GPT-4o under the same setup, and it produces consistent results, confirming that the merit-guided prompts generated by lightweight LLMs outperform the Basic Prompt in quality.

These studies collectively show that lightweight LLMs struggle with vague or complex optimization instructions but can perform competitively when provided with structured, merit-aligned guidance.

B Identifying Effective Prompt Merits via Comparative Evaluation

To systematically analyze effective prompt characteristics, we conduct comparative evaluations. All experiments are implemented using the lightweight LLM M_{ref} (Qwen2.5-7B-Instruct), which serves as: (1) the merit-guided prompt optimization model, (2) the inference model, and (3) the prompt rewriting model. Below are the details for each evaluation step.

Prompt Rewriting. We begin by using a rewriting prompt (shown in Fig. 6) to generate multiple reformulations of each raw question. Specifically, five rewritten versions are produced for each input prompt to preserve diversity. Below is an example from the Alpaca dataset, where the raw prompt results in an incorrect response, while Rewrite 1 and Rewrite 2 yield correct answers under the same inference model:

Raw: What does DNA stand for?

Rewrite 1: Can you tell me what the acronym DNA represents?

Rewrite 2: What is the full form of DNA?

Prompt Effectiveness Evaluation. To evaluate prompt quality at the prompt level (i.e., independent of response content), we design a comparison prompt that asks DeepSeek-R1 to assess which

of the two prompt variants is better and explain why, as illustrated in Fig. 8. The summarized features provided by DeepSeek-R1 are then collected for merit discovery. Below presents two examples from DeepSeek’s evaluations, highlighting which prompt was preferred and why, along with the associated merit types.

Example 1

DeepSeek Evaluation:

###Explicit Guidance: The Golden Prompt provides a clearer direction by specifying the focus on "the importance of customer service," which helps the responder structure their thoughts around a central theme rather than a vague topic.

###Contextual Richness: By emphasizing "the importance," the Golden Prompt encourages the responder to include reasons, examples, or impacts, leading to a more substantive and engaging tweet.

###Engagement Potential: The Golden Prompt is more likely to generate a tweet that resonates with readers, as discussing "importance" naturally invites the responder to highlight value or benefits, making the content more compelling and shareable.

In contrast, the Sliver Prompt ("Craft a tweet regarding customer service experiences") is too broad and could lead to generic or unfocused responses, such as simple complaints or praises without deeper insight.

Example 2

DeepSeek Evaluation:

###Precision in Terminology: The Golden Prompt uses the more accurate linguistic term "form" instead of "conjugation," which is more appropriate for discussing verb structures in English grammar.

###Clarity of Focus: By placing the verb phrase in quotation marks ("will have helped"), it clearly identifies the exact linguistic element to be analyzed, reducing ambiguity.

###Technical Accuracy: The phrasing aligns better with grammatical analysis,

as "form" encompasses tense, aspect, and mood, which are all relevant to "will have helped" (future perfect tense).

###Conciseness: It removes the unnecessary framing of "Identify the conjugation of the verb:" and gets straight to the point, making it more efficient for analysis.

Prompt Preference Judgement. We use a prompt template (see Fig. 9) to conduct direct pairwise comparisons between two prompts, with judgments provided by DeepSeek-R1. The positions of the prompts are randomly assigned to mitigate positional bias.

Response Scoring. The quality of the model’s response to each prompt is evaluated using DeepSeek-R1. The scoring template is displayed in Fig. 13, inspired by Yang et al. (2025).

Prompt Optimization. The prompt template used to refine raw prompts into merit-guided optimized versions is shown in Fig. 14.

C POP Dataset Construction and Training Details

In dataset construction and prompt optimizer training, the lightweight LLM M_{ref} (Qwen2.5-7B-Instruct) serves as: (1) the prompt degradation model, (2) the inference model, (3) the merit-guided prompt optimization model, and (4) the base model for prompt optimization learning.

Prompt Degradation. To simulate noisy user inputs—such as “*how can i go fr sigapor?*”—we degrade 10% of the raw prompts from the Alpaca and BPO datasets using a base model. The prompt used for degradation is shown in Fig. 7. Below is an example from the Alpaca dataset:

Raw: Describe the atmosphere at the beach.

Degraded: describ the atmossphre at the bheach.

DPO Input Construction. To train MePO under the DPO objective, we formulate each training instance using the format shown in Fig. 10, where S_P is the silver (unoptimized) prompt P_{silver} , S_R is its corresponding response R_{silver} , and G_R is the preferred response R_{golden} .

Training Configuration. We fine-tune MePO based on Qwen2.5-7B-Instruct² using 4 NVIDIA RTX A5000 GPUs. Each input instruction x is truncated or padded to 2000 tokens. Training hyperparameters and cost are summarized in Table 6.

Qwen2.5-7b-instruct	
Node1	1
Per_GPU_batch	1
Accumulations	4
HyperParams	Epoch=2, Seq Len=2048, lr=1e-6, beta=0.01, Top P=0.95, Temperature=0.8, loss type='sigmoid',
Train (2 epoch)	1.5 day

Table 6: Training setup and hyperparameters used for fine-tuning MePO.

D Dataset Information and Experimental Implementation

We evaluate MePO on four multiple-choice and one generative benchmark dataset:

- **ARC⁷**: A grade-school science QA benchmark divided into ARC-Easy (2,376 questions) and ARC-Challenge (1,172 questions).
- **GSM8K⁸**: A dataset of 1,319 grade-school math word problems requiring free-form answer generation.
- **BBH⁹**: A suite of 25 complex reasoning tasks from BigBench.
- **PiQA¹⁰**: A multiple-choice dataset assessing commonsense physical reasoning, with 1,838 questions.

Evaluation Protocol. GSM8K, PiQA, and BBH are evaluated in a 3-shot setting to facilitate answer extraction, while all other datasets are evaluated in a zero-shot setting. For consistency, all BBH tasks are reformatted into multiple-choice format, except for `multistep_arithmetic_two`, `object_counting`, and `word_sorting`, which retain their original formats. PiQA, which includes two-solution questions, is also converted to multiple-choice format.

Unlike previous Qwen and LLaMA evaluations (Yang et al., 2024; Grattafiori et al., 2024), which select the option (e.g., A, B, or 1, 2) based on the highest next-token logit probability, we extract the answer from the model’s generated reason-

⁷https://huggingface.co/datasets/allenai/ai2_arc

⁸<https://huggingface.co/datasets/openai/gsm8k>

⁹<https://huggingface.co/datasets/lucaemon/bbh/>

¹⁰<https://huggingface.co/datasets/ybisk/piqa>

Inference Model	Optimizer	ARC-Easy	ARC-Challenge	GSM8K	BBH	PiQA	Avg.
LLaMA3.2-1b-instruct	-	27.02	23.13	16.38	31.05	49.56	29.43
	Inference Model	28.79	24.23	20.24	32.29	49.73	31.06
	BPO	27.99	24.04	15.85	31.77	49.89	29.91
	FIPO	25.19	23.55	21.30	31.72	50.27	30.41
	MePO	29.42	24.40	21.46	33.40	50.44	31.82
LLaMA3.2-3b-instruct	-	73.32	55.46	63.23	37.55	66.32	59.18
	Inference Model	78.16	60.75	63.00	40.12	69.37	62.28
	BPO	74.75	55.74	56.71	40.01	67.14	58.87
	FIPO	78.62	60.98	65.13	40.27	69.53	62.91
	MePO	79.55	61.43	65.66	43.56	70.29	64.10

Table 7: Downward compatibility: Response evaluation of optimizer in lightweight inference models.

ing. Specifically, we use format-specific instructions—e.g., for zero-shot ARC: “Reply with the answer option starting with ##, like ##A, ##B, ##C, or ##D”—and extract the answer that follows the ## marker in the response. This approach better reflects the model’s reasoning ability during answer selection.

Prompt Templates. Fig. 12 illustrates the prompt formats used for GSM8K, PiQA, and ARC. For BBH, multiple-choice tasks follow the PiQA format, with in-context examples prepended to the test query along with the corresponding golden answers. Math-related tasks in BBH adopt the GSM8K format.

Baseline Implementations. We compare MePO with two state-of-the-art prompt optimizers. The prompt instructions used for inference models serving as prompt optimizers are shown in Fig. 11. Details of the baselines are as follows:

- **FIPO** is trained on Tulu2-70B-DPO¹¹, using a GPT-3.5/GPT-4-generated preference dataset focused on chain-of-thought reasoning.
- **BPO** is trained on LLaMA2-7B¹², with a dataset built from human-preferred prompts optimized by GPT-3.5.

All optimization models, including MePO, are loaded in 8-bit precision for prompt generation. Inference models are run in 4-bit mode with a generation length below 512 to reduce memory overhead.

D.1 Case Study: Downward Compatibility

To further analyze the downward compatibility of optimization models, we evaluate performance us-

ing two smaller inference models: LLaMA-3.2-3B-Instruct¹³ and LLaMA-3.2-1B-Instruct¹⁴.

As shown in Table 7, both BPO and FIPO exhibit performance degradation when paired with lightweight inference models. For instance, under LLaMA-3.2-3B-Instruct, compared to using the inference model itself as the optimizer, BPO drops by 3.41% on ARC-Easy, 5.01% on ARC-Challenge, 6.29% on GSM8K, and 2.23% on PiQA. Under LLaMA-3.2-1B-Instruct, FIPO drops by 3.60% on ARC-Easy, and BPO drops by 4.39% on GSM8K. Both also show decreased performance on BBH across the two lightweight inference models. In contrast, MePO consistently improves performance across all datasets, demonstrating strong downward compatibility and robustness in low-resource settings.

D.2 Case Study: Online-Based Discrete APO Comparison

To further validate MePO’s effectiveness in optimizing prompts or instructions with refined meta-prompts, we compare MePO with three discrete APO methods on GSM8K:

- **APO**, a prompt optimization method that performs optimization via textual gradients using GPT-3.5-Turbo-Instruct in a text-based dialogue setting;
- **Iterative APE**, a template-based strategy in which GPT-3.5 generates a pool of candidate prompts, followed by selection based on development set performance;
- **PE2**, a prompt engineering method that optimizes prompts through three key components using GPT-3.5-Turbo-Instruct.

¹¹<https://huggingface.co/Junrulu/FIPO-IPL-IP0-Tulu2-70B>

¹²<https://huggingface.co/THUDM/BPO>

¹³<https://huggingface.co/meta-llama/Llama-3.2-3B-Instruct>

¹⁴<https://huggingface.co/meta-llama/Llama-3.2-1B-Instruct>

The meta-prompts for each method are listed in Table 8, adapted from Ye et al. (2024), and are prepended before the ##Answer segment shown in Fig. 12(a).

Optimizer	Meta Prompt
APO	Given the scenario, perform necessary calculations and provide a step-by-step explanation to arrive at the correct numerical answer. Consider all information provided.
Iterative APE	Let’s dissect this and tackle it gradually, one phase at a time.
PE2	Let’s solve the problem step-by-step and calculate the required total value correctly.

Table 8: GSM8K meta-prompts used for each baseline discrete APO method.

D.3 BBH Experimental Results

Detailed results for the 23 BBH tasks are presented in the following tables: Table 9 (raw prompts), Table 10 (optimized by the inference model), Table 11 (training-free merit-guided optimization), Table 12 (BPO), Table 13 (FIPO), and Table 14 (MePO).

-	Qwen2-7b	Tulu2-7b -dpo	LLaMA2-7b -chat-hf	LLaMA3-8b -instruct	LLaMA2-13b -chat-hf	LLaMA3.2-3b -instruct	LLaMA3.2-1b -instruct	Tulu2-70b -dpo	LLaMA3.3-70b -instruct
date_understanding	44.8	30.8	28.8	31.6	42	25.6	27.2	62	60.8
disambiguation_qa	63.2	32.8	31.2	32.4	30	36.4	33.6	68.8	65.6
hyperbaton	75.6	48.4	48.4	48.4	48.4	51.6	48.4	78	87.6
logical_deduction_five_objects	49.2	22.4	18	28	33.2	24.4	19.6	50	58.8
logical_deduction_seven_objects	42.4	17.6	14.8	30	27.2	30.4	15.2	42.8	54.4
logical_deduction_three_objects	66.4	32.8	31.6	33.6	36.4	35.2	31.6	71.2	75.6
movie_recommendation	48	26.8	22.4	33.2	23.6	23.2	22.4	52.8	76.4
penguins_in_a_table	54.11	21.23	22.6	31.51	46.58	26.71	23.97	50.68	67.12
reasoning_about_colored_objects	47.6	32.4	32	16.4	29.2	30	16.4	46	51.6
ruin_names	37.2	24	30.8	35.2	27.2	41.6	28.4	66.8	51.2
salient_translation_error_detection	32	26.8	23.2	21.6	28	34	16.8	52	71.6
snarks	60.11	44.94	46.07	46.07	50.56	57.3	46.07	81.46	74.16
word_sorting	88	82.4	80	84	80	52	1.2	100	100
temporal_sequences	61.2	19.2	28.4	28.4	33.2	28	28.4	62	83.6
tracking_shuffled_objects_five_objects	28.4	19.2	20	20.8	30.8	20	20	17.2	40.4
tracking_shuffled_objects_seven_objects	13.2	24	14.4	13.6	22	11.2	12.8	13.6	18.4
tracking_shuffled_objects_three_objects	27.2	29.2	24.4	26.8	49.2	30	30.8	28	31.2
causal_judgement	62.57	49.2	51.87	51.34	49.2	47.59	51.87	53.48	79.68
formal_fallacies	54	52.8	53.2	53.6	52.4	51.6	53.2	52	77.2
navigate	57.6	53.6	42	41.6	41.6	41.2	42	52.8	57.2
web_of_lies	46.4	48.8	48.8	49.2	54	44.4	46.4	46.8	68.4
sports_understanding	70	50	46.4	70.4	54	53.6	43.2	45.6	66.8
boolean_expressions	34.8	56.8	52.4	54.8	51.2	48	54	30.4	62.8
multistep_arithmetic_two	33.6	12.4	9.6	36	14	46	16	16.8	85.6
object_counting	36	33.2	32.8	50.4	33.2	48.8	46.8	49.6	90

Table 9: Results of raw BBH dataset across the evaluated inference models.

-	Qwen2-7b	Tulu2-7b -dpo	LLaMA2-7b -chat-hf	LLaMA3-8b -instruct	LLaMA3.2-3b -instruct	LLaMA3.2-1b -instruct
date_understanding	54.4	34	34	32	27.6	31.2
disambiguation_qa	62.8	34	32.4	34.4	40.4	36.8
hyperbaton	77.6	49.2	49.2	49.2	51.6	48.4
logical_deduction_five_objects	50.4	22.8	19.2	36.4	29.2	20.8
logical_deduction_seven_objects	44	17.2	15.6	31.2	29.6	16.4
logical_deduction_three_objects	74.8	34.8	32	38.8	40	32
movie_recommendation	51.6	30.4	23.6	35.2	24.4	26.8
penguins_in_a_table	56.16	26.03	23.29	32.19	28.77	23.29
reasoning_about_colored_objects	48.4	35.2	34	18.8	28.4	18
ruin_names	38.4	26.4	32.4	38	43.6	28.4
salient_translation_error_detection	38	21.2	24.8	23.6	35.2	20.4
snarks	65.73	48.88	49.44	47.19	57.3	46.07
word_sorting	90.8	84	84	86.4	60	1.2
temporal_sequences	62.8	31.2	33.2	30.4	28.4	28.4
tracking_shuffled_objects_five_objects	29.6	29.6	30.8	32.4	20	22.4
tracking_shuffled_objects_seven_objects	15.2	25.6	23.2	14.4	14	13.6
tracking_shuffled_objects_three_objects	28	30	26	28.4	31.2	32
causal_judgement	72.73	52.41	53.2	50.8	55.61	52.41
formal_fallacies	54.8	54.4	53.2	51.6	52	53.2
navigate	60.8	56	43.2	42.8	41.6	42
web_of_lies	47.2	48.8	48.8	50	48.8	47.6
sports_understanding	75.6	63.45	48	72	57.6	45.6
boolean_expressions	41.6	57.2	53.6	52.4	47.6	54
multistep_arithmetic_two	36.4	10.8	12	37.2	54	18
object_counting	39.2	33.2	37.2	58.4	56	48.4

Table 10: Results of the BBH dataset optimized by the inference model across the evaluated inference models.

-	Qwen2-7b	Tulu2-7b-dpo	LLaMA2-7b-chat-hf	LLaMA3-8b-instruct
date_understanding	55.6	36	34	48
disambiguation_qa	66	37.2	35.2	35.2
hyperbaton	78.4	50	48.4	49.6
logical_deduction_five_objects	55.6	25.2	19.2	34.8
logical_deduction_seven_objects	44	28.4	21.6	38.8
logical_deduction_three_objects	75.2	34.4	42.4	40.4
movie_recommendation	53.2	31.2	24.4	37.6
penguins_in_a_table	56.16	27.4	23.29	36.99
reasoning_about_colored_objects	50.8	34	35.2	20.8
ruin_names	38	28.4	40	39.6
salient_translation_error_detection	38.4	34	37.2	18.8
snarks	62.36	50	50.56	48.88
word_sorting	94.4	87.2	87.6	90
temporal_sequences	64.4	36.8	36.8	31.2
tracking_shuffled_objects_five_objects	29.6	38.4	34.8	33.2
tracking_shuffled_objects_seven_objects	14.4	24.8	23.2	14.4
tracking_shuffled_objects_three_objects	30	31.6	31.6	31.2
causal_judgement	55.08	52.41	51.87	52.41
formal_fallacies	54	54	54	52
navigate	57.2	58.8	43.2	42.4
web_of_lies	52.4	49.6	49.2	51.6
sports_understanding	76	64.8	49.2	79.6
boolean_expressions	41.6	60.4	53.6	50
multistep_arithmetic_two	41.2	12.6	12	41.6
object_counting	39.2	36.4	58.4	50.8

Table 11: Results of the BBH dataset optimized by training-free merit-guidance across the evaluated inference models.

-	Qwen2-7b	Tulu2-7b-dpo	LLaMA2-7b-chat-hf	LLaMA3-8b-instruct	LLaMA3.2-3b-instruct	LLaMA3.2-1b-instruct
date_understanding	56.8	33.2	35.2	48.8	29.6	29.2
disambiguation_qa	63.2	34	34.4	40.4	44.4	34.4
hyperbaton	80	53.2	50	49.6	51.6	48.8
logical_deduction_five_objects	52.8	25.2	24	31.2	29.2	20.4
logical_deduction_seven_objects	46.4	28.8	19.2	35.6	24.8	15.2
logical_deduction_three_objects	78.4	32.8	45.6	39.2	34.8	32
movie_recommendation	61.2	22.8	24.4	37.6	23.2	26.8
penguins_in_a_table	54.11	23.97	23.29	36.3	29.45	23.97
reasoning_about_colored_objects	52.4	34	35.6	20.8	28.4	18
ruin_names	38.4	28	36.8	35.6	38	28.4
salient_translation_error_detection	38.4	27.6	35.2	21.2	33.6	18.8
snarks	66.85	53.37	53.37	48.88	53.93	46.07
word_sorting	92.4	90	87.6	90	72	3.2
temporal_sequences	66.4	26.4	38.4	30.8	35.6	28.4
tracking_shuffled_objects_five_objects	29.6	29.2	31.2	34	20	19.6
tracking_shuffled_objects_seven_objects	14.4	28	23.2	15.2	14	13.2
tracking_shuffled_objects_three_objects	27.6	31.6	31.6	31.6	31.6	32
causal_judgement	53.48	50.8	52.04	52.94	52.94	51.87
formal_fallacies	52	53.2	54.14	54	52.4	53.2
navigate	59.6	58.4	44.8	48.4	44.8	42
web_of_lies	49.2	51.6	48.8	49.6	49.2	48.8
sports_understanding	76.4	61.2	53.2	74.8	55.2	45.6
boolean_expressions	45.6	63.6	55.6	53.2	50.4	54
multistep_arithmetic_two	36.8	10	17.6	37.2	46.8	16.8
object_counting	34.4	32.4	34.8	50	54.4	43.6

Table 12: Results of the BBH dataset optimized by BPO across the evaluated inference models.

-	Qwen2-7b	Tulu2-7b -dpo	LLaMA2-7b -chat-hf	LLaMA3-8b -instruct	LLaMA3.2-3b -instruct	LLaMA3.2-1b -instruct
date_understanding	55.2	38.4	29.2	36	30	27.2
disambiguation_qa	61.6	34.8	31.2	40.4	32	33.2
hyperbaton	78.8	33.2	48.4	60	52	48.4
logical_deduction_five_objects	50	21.6	22.8	34.4	23.6	20.4
logical_deduction_seven_objects	46.4	25.2	20.4	39.2	23.2	15.6
logical_deduction_three_objects	70.8	30.4	42	39.2	38.8	32
movie_recommendation	62	36.8	23.97	31.6	34.4	24.4
penguins_in_a_table	54.11	36.3	23.29	30.82	52.05	23.29
reasoning_about_colored_objects	50.8	31.6	33.2	24.8	25.6	16.4
ruin_names	46	34.4	36	38.8	40	29.2
salient_translation_error_detection	37.6	32.4	33.6	24.4	37.2	18.8
snarks	62.92	46.07	50.56	46.07	55.62	46.07
word_sorting	92.4	92.4	86.4	92	70	4.8
temporal_sequences	63.2	27.6	28.4	28.8	40	28.8
tracking_shuffled_objects_five_objects	31.2	24.8	32.4	36.8	20.8	20.8
tracking_shuffled_objects_seven_objects	16.8	25.6	24.8	14.4	14.4	16.4
tracking_shuffled_objects_three_objects	31.2	30.4	31.6	30.4	31.6	30.8
causal_judgement	62.57	53.48	51.87	51.34	55.08	52.41
formal_fallacies	53.2	54.4	54	52.4	53.2	53.2
navigate	59.2	60.8	42	46.4	53.2	42
web_of_lies	58.8	48	53.2	45.2	49.2	48.8
sports_understanding	81.2	66.4	47.26	71.6	55.6	46
boolean_expressions	39.6	58.8	53.6	61.2	39.6	54.4
multistep_arithmetic_two	16.4	11.6	17.2	39.6	24.4	16.4
object_counting	36.4	33.6	61.2	50.8	55.2	43.2

Table 13: Results of the BBH dataset optimized by FIPO across the evaluated inference models.

-	Qwen2-7b	Tulu2-7b -dpo	LLaMA2-7b -chat-hf	LLaMA3-8b -instruct	LLaMA2-13b -chat-hf	LLaMA3.2-3b -instruct	LLaMA3.2-1b -instruct	Tulu2-70b -dpo	LLaMA3.3-70b -instruct
date_understanding	57.77	36.4	35.6	49.2	52.4	34	31.2	50	64.8
disambiguation_qa	66.4	38.4	37.2	48.8	54.8	44.8	40.8	68.8	71.6
hyperbaton	81.6	54.4	50.4	50.4	54.4	56.8	48.4	62	90.8
logical_deduction_five_objects	58.8	27.2	24.4	47.2	42.4	30	23.6	52	66.8
logical_deduction_seven_objects	49.2	29.6	20.4	39.6	36.8	21.2	18.4	51.6	63.2
logical_deduction_three_objects	82	36.8	51.2	42.8	39.2	40.4	34.8	65.6	80.8
movie_recommendation	63.2	38.4	26.8	38.4	29.2	32	27.2	34	77.6
penguins_in_a_table	56.85	28.77	25.34	38.36	54.79	31.51	25.34	51.37	71.23
reasoning_about_colored_objects	56.8	35.6	38	23.6	41.6	26	18.4	58.8	63.2
ruin_names	38	27.6	41.6	42	28.8	45.6	28.4	73.6	59.2
salient_translation_error_detection	43.2	36.8	38.8	28.8	31.2	39.2	23.6	54.8	72
snarks	67.42	53.37	51.12	49.44	66.85	60.11	46.07	83.15	76.97
word_sorting	95.2	92.4	90	92	86.4	84	3.2	100	100
temporal_sequences	38	37.6	39.2	33.2	42	29.2	28.8	66.8	84.4
tracking_shuffled_objects_five_objects	31.2	39.6	33.2	35.2	32.8	20.4	23.2	21.2	28
tracking_shuffled_objects_seven_objects	15.2	30	26.4	16.4	28.4	18.8	15.2	18	30.4
tracking_shuffled_objects_three_objects	33.6	32.4	33.2	33.2	64.8	31.6	31.6	29.6	31.6
causal_judgement	49.2	56.68	51.87	51.79	56.15	62.57	52.41	59.36	80.21
formal_fallacies	51.2	57.2	54.8	53.2	52.8	53.2	53.2	59.2	83.2
navigate	61.6	58.8	44.4	44.8	48.4	47.2	42	60	42.4
web_of_lies	51.2	55.2	51.6	51.6	59.2	51.2	49.6	49.6	71.6
sports_understanding	77.6	67.2	54.8	77.2	61.2	58	46	49.2	75.6
boolean_expressions	48.4	61.2	54	54.4	54	52.8	55.6	36.4	66.4
multistep_arithmetic_two	45.6	14.4	14.4	41.2	16	57.6	18.4	18	84.4
object_counting	39.6	35.2	60.4	55.2	34	60.8	49.6	53.6	86.8

Table 14: Results of the BBH dataset optimized by MePO across the evaluated inference models.

E Prompt Templates

The prompts used in this work are listed below:

Prompt for EvoPrompt Comparison

You are an expert in prompt evaluation. Given two prompts derived from the same **##Basic Prompt##**—**##Prompt 1##** and **##Prompt 2##**—determine which one is better overall for eliciting high-quality responses from a language model and information related to **##Basic Prompt##**.

##Basic Prompt##:
B_P

##Prompt 1##:
S_P

##Prompt 2##:
G_P

Which is better? Please respond with only `1` or `2`, followed by a brief explanation if necessary.

Figure 5: Prompt used to evaluate EvoPrompt Prompt and Optimal Prompts generated by the EvoPrompt algorithm under lightweight LLMs.

Prompt for Question Rewrite

Given the following sentences, generate five more sentences that express the same meaning but use different words.
Original sentences:
{}

Generate five alternative versions:

Figure 6: Prompt used to rewrite raw questions.

Prompt for Degrade Prompt

Your task is to destroy a **##Good Prompt##** to make it significantly less effective.

Your output should retain the general topic but degrade the clarity, grammar, usefulness, and precision of the **##Good Prompt##**.

##Good Prompt##:
S_P

Only give me the content of **##Bad Prompt##** in English, do not contain any other information and Chinese (e.g., any postfix like 'Bad Prompt', etc.).

##Bad Prompt##:

Figure 7: Prompt used to generate degraded prompts. S_P denotes the original raw prompt.

Prompt for Optimal Pattern Evaluation

You are an expert prompt evaluator. Given two prompts: **##Sliver Prompt##** and **##Golden Prompt##**, both derived from the same raw input, and knowing that the Golden Prompt consistently leads to better responses, please explain in English why the Golden Prompt is better.

##Sliver Prompt##:
S_P

##Golden Prompt##:
G_P

Comments:

Figure 8: Prompt used to evaluate the effectiveness of two prompts. S_P denotes the prompt yielding a lower-scoring response; G_P yields a higher-scoring response.

Prompt for Prompt Comparison

You are an expert in prompt evaluation. Given two prompts derived from the same original input—**##Prompt 1##** and **##Prompt 2##**—determine which one is better overall for eliciting high-quality responses from a language model.

##Prompt 1##:
S_P

##Prompt 2##:
G_P

Which is better? Please respond with only `1` or `2`, followed by a brief explanation if necessary.

Figure 9: Prompt used to compare raw and optimized prompts. The two prompts are randomly placed in S_P and G_P to mitigate position bias.

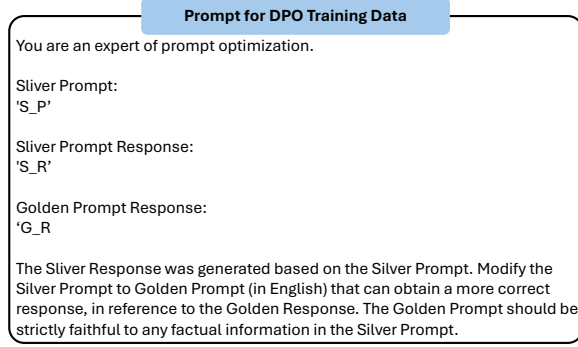


Figure 10: Prompt used to construct DPO training inputs. S_P , S_R , and G_R denote P_{silver} , R_{silver} , and R_{golden} , respectively.

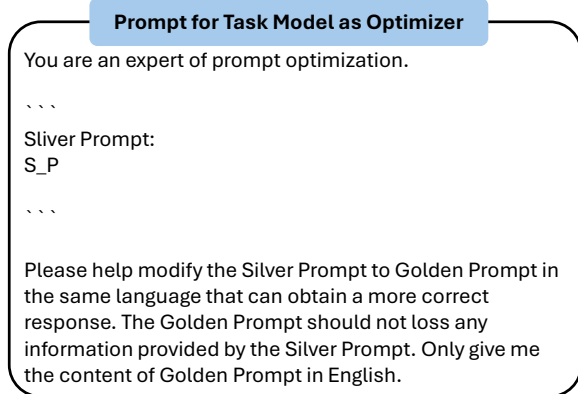
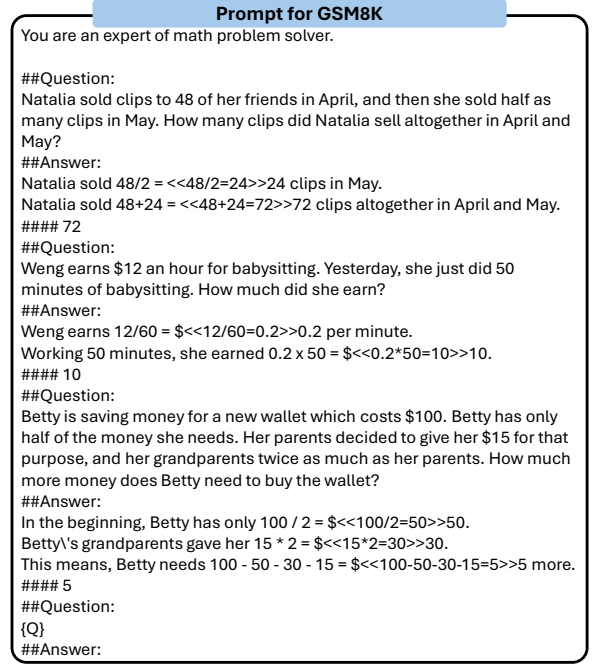
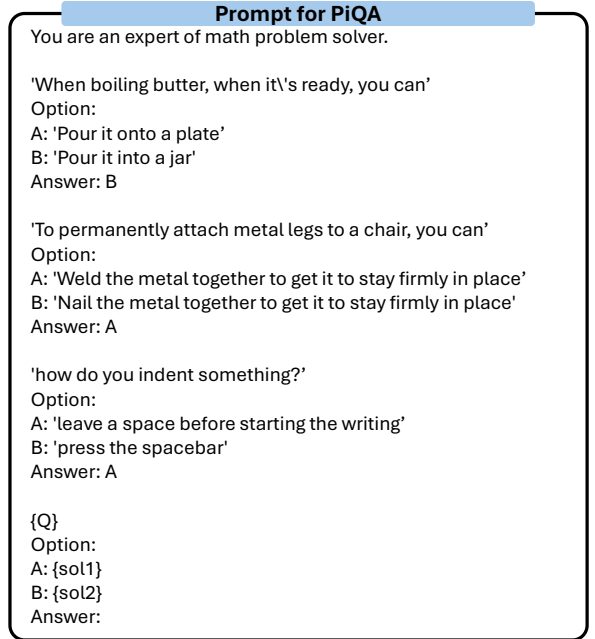



Figure 11: Prompt used to optimize prompts for inference model. S_P denotes the original raw prompt.



(a) GSM8K, BBH math-related tasks



(b) PiQA



(c) ARC, BBH multiple-choice tasks

Figure 12: Prompt formats used for downstream task evaluation.

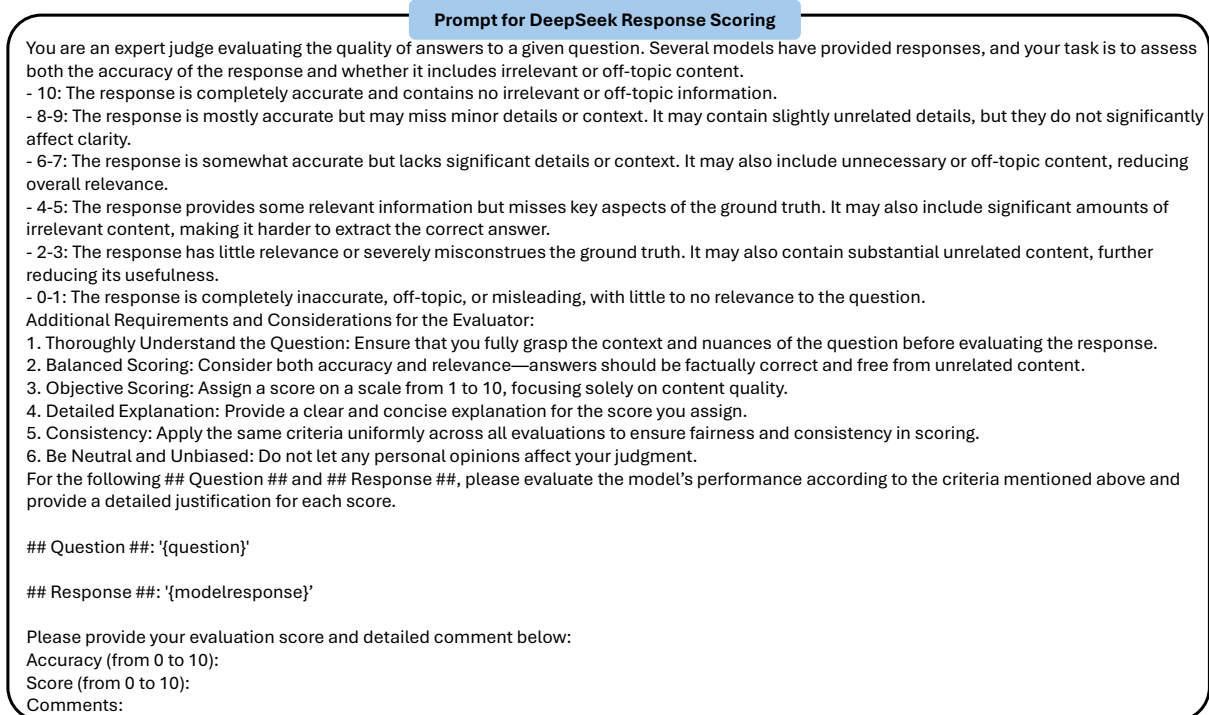


Figure 13: Prompt used by DeepSeek-R1 to score responses. {question} is replaced by the prompt; {modelresponse} is replaced by the model's response.

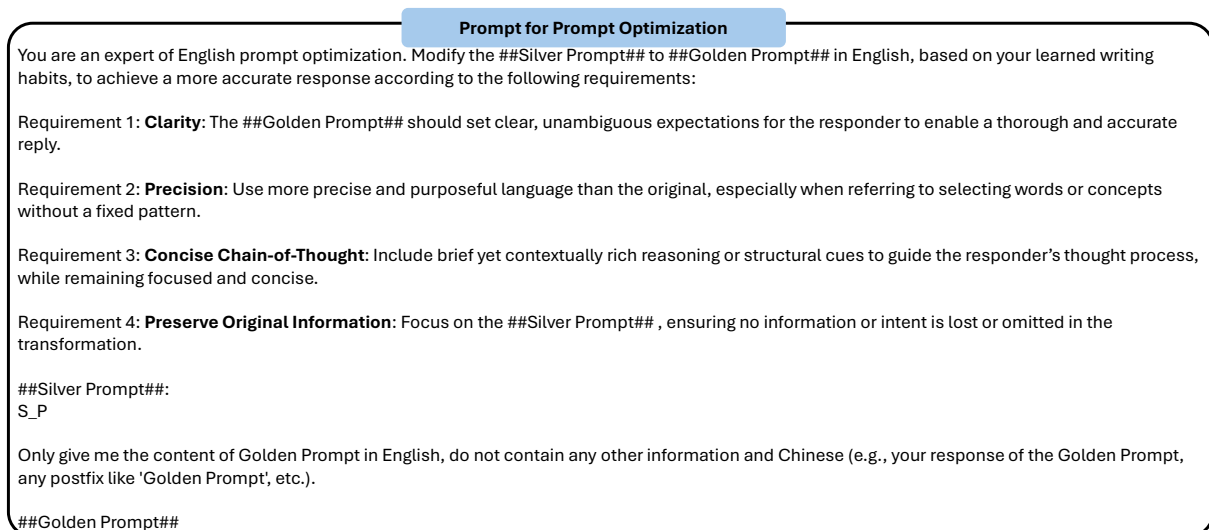


Figure 14: Prompt for Prompt Optimization. S_P denotes the prompt that needs to be optimized.