

LLM Prompt Duel Optimizer: Efficient Label-Free Prompt Optimization

Yuanchen Wu^{1*†} Saurabh Verma² Justin Lee² Fangzhou Xiong²
Poppy Zhang² Amel Awadelkarim² Xu Chen² Yubai Yuan¹ Shawndra Hill²
¹Department of Statistics, The Pennsylvania State University ²Meta

Abstract

Large language models (LLMs) are highly sensitive to prompts, but most automatic prompt optimization (APO) methods assume access to ground-truth references (e.g., labeled validation data) that are costly to obtain. We propose the **Prompt Duel Optimizer (PDO)**, a sample-efficient framework for *label-free* prompt optimization based on pairwise preference feedback from an LLM judge. PDO casts prompt selection as a dueling-bandit problem and combines (i) *Double Thompson Sampling* to prioritize informative comparisons under a fixed judge budget, with (ii) *top-performer guided mutation* to expand the candidate pool while pruning weak prompts. Experiments on BIG-bench Hard (BBH) and MS MARCO show that PDO consistently identifies stronger prompts than label-free baselines, while offering favorable quality–cost trade-offs under constrained comparison budgets.

1 Introduction

LLM performance hinges on well-crafted prompts that unlock their capabilities (Sun et al., 2023). Creating effective prompts typically requires extensive trial-and-error or task-specific techniques (e.g., chain-of-thought prompting for reasoning tasks; Wei et al., 2022), which often do not transfer across tasks or domains. This limitation motivates Automatic Prompt Optimization (APO) (Ramnath et al., 2025), which iteratively generates and evaluates candidate prompts to discover high-performing instructions.

Despite encouraging results across diverse tasks, most APO methods (Zhou et al., 2022; Yang et al., 2024; Fernando et al., 2023; Pryzant et al., 2023) rely on **reference-based supervision** (e.g., ground-truth labels, gold references, or label-derived rewards) to score candidates on validation sets. In

many practical settings, however, obtaining such supervision at scale is costly and slow (Ratner et al., 2017). Still, prompt optimization remains essential; for example, in industrial LLM-based text classification, practitioners need reasonably good prompts to initiate deployment before large human-labeled datasets are available (Wagner, 2024). Such challenges raise a central question:

- *Can we optimize prompts without ground-truth label references?*

One direction is to reduce reliance on human annotation by adopting automatic evaluation of model outputs. Recent work suggests that LLMs can serve as judges of model outputs, including in reference-free settings (Zheng et al., 2023; Liu et al., 2023; Gu et al., 2024). As prompt quality is reflected in generated responses, evaluating prompts with an LLM judge is a natural approach. In this setting, rather than scoring each output independently, it is often preferable to rely on **pairwise preference feedback**: a judge compares outputs from two prompts on the same input and selects the preferred one. Pairwise comparisons typically yield a more reliable signal than direct pointwise scoring, which is prone to calibration issues (Liu et al., 2024).

LLM preference feedback introduces two challenges. *First*, LLM judges are noisy: calls are non-deterministic (He and Lab, 2025), judgments can exhibit position and verbosity biases (Zheng et al., 2023), and task complexity may amplify these effects. *Second*, preference-based evaluation scales quadratically with the number of candidates; each comparison requires an LLM API call and incurs monetary cost, making exhaustive evaluation infeasible.

To address these issues, we introduce the **Prompt Duel Optimizer (PDO)**, which casts label-free prompt optimization as a dueling-bandit problem where the cost is the number of judge compar-

*Work done during an internship at Meta.

†Code is available at <https://github.com/meta-llama/prompt-ops>

isons. PDO combines *Double Thompson Sampling (D-TS)* to focus comparisons on informative prompt pairs with a top-performer *mutate-and-prune* strategy that expands the pool while discarding weak candidates. We further assess judge reliability via cross-family re-evaluation to probe circularity and show that PDO’s gains are robust across judge choices.

We summarize our contributions as follows:

1. We formalize label-free prompt optimization with LLM-judged pairwise comparisons as a dueling-bandit problem.
2. We propose PDO, integrating D-TS for sample-efficient prompt selection with top-performer guided mutation that steers search toward stronger regions and yields better candidates.
3. We show empirically on BBH and MS MARCO that PDO outperforms label-free baselines under fixed judge budgets, and analyze quality–cost trade-offs via budget and runtime comparisons.
4. We analyze judge noise and potential circularity, and demonstrate that PDO’s gains persist under alternative judge families.

2 Preliminaries and Problem Setup

In this section, we establish the connection between preference-based prompt optimization and dueling bandits.

2.1 Background: Dueling Bandits

In the K -armed dueling bandit problem (Bengs et al., 2021), at each round the decision-maker selects two arms i and j to duel, and the comparison yields a stochastic preference outcome. Specifically, the probability that i is preferred to j is

$$\mu(i, j) = \Pr(i \succ j).$$

A *Condorcet winner* is an arm i^* such that $\mu(i^*, j) > 0.5$ for all $j \neq i^*$. When no Condorcet winner exists, a standard choice rule is the Copeland criterion: the *Copeland score* of arm i counts the number of opponents it beats with probability greater than $1/2$:

$$C(i) = |\{j \neq i : \mu(i, j) > \frac{1}{2}\}|.$$

A *Copeland winner* is then an arm $i^* \in \arg \max_i C(i)$.

2.2 Prompt Optimization Through the Lens of Dueling Bandits

Building on these definitions, we formalize preference-based prompt optimization in the dueling bandits setting.

Problem Setup. Let \mathcal{X} denote the input space and $\mathcal{P} = \{p_1, \dots, p_K\}$ a finite set of candidate prompts. Let $D_{\mathcal{X}}$ be a distribution over \mathcal{X} . For $p \in \mathcal{P}$ and $x \sim D_{\mathcal{X}}$, let $f_p(x)$ denote the LLM output when applying prompt p to input x . To compare the two prompts $p_i, p_j \in \mathcal{P}$ on the same input x , we query an LLM judge and record a binary preference:

$$\text{Judge}_x(p_i, p_j) = \begin{cases} 1, & \text{if } f_{p_i}(x) \succ f_{p_j}(x), \\ 0, & \text{otherwise.} \end{cases}$$

Given a set of unlabeled examples $\{x_i\}_{i=1}^n$, the empirical estimate of $\mu(p_i, p_j)$ becomes

$$\hat{\mu}(p_i, p_j) = \frac{1}{n} \sum_{k=1}^n \mathbf{1}[\text{Judge}_{x_k}(p_i, p_j) = 1].$$

Prompt Optimization Objective. The goal of prompt optimization is to identify a prompt that maximizes task performance. In the absence of ground-truth references, we use pairwise preferences as a practical proxy for selecting high-quality prompts. Using empirical estimates $\hat{\mu}(p_i, p_j)$, we select the Condorcet winner when it exists, or otherwise the Copeland winner. In practice, limited API budgets make exhaustive estimation ($O(n|\mathcal{P}|^2)$ comparisons) infeasible, motivating sample-efficient methods that adaptively target informative pairs while aligning with the Condorcet/Copeland objectives.

3 Prompt Duel Optimizer (PDO)

To address the above challenge, we propose the **Prompt Duel Optimizer (PDO)**, an algorithm designed to identify high-performing prompts under limited comparison budgets. PDO combines two complementary components: (1) *Double Thompson Sampling (D-TS)* (Wu and Liu, 2016), a Bayesian strategy for efficient pairwise evaluation that targets Copeland-optimal winners; and (2) *Top-Performer Guided Mutation*, which adaptively mutates top-performing prompts to expand coverage of the search space.

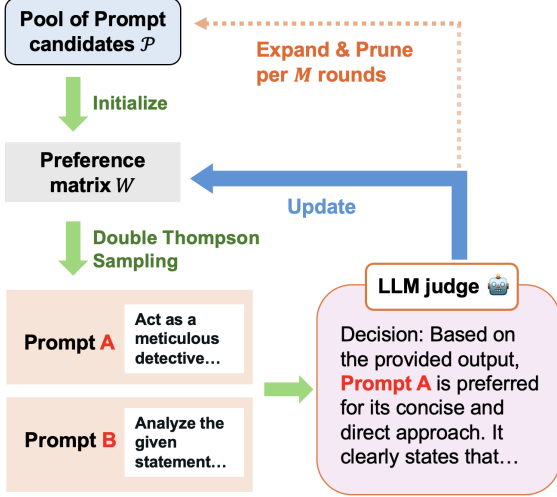


Figure 1: Workflow of the Prompt Duel Optimizer.

3.1 Efficient Prompt Selection via Double Thompson Sampling

D-TS extends Thompson Sampling (Agrawal and Goyal, 2012) to the dueling bandit setting, where feedback comes from pairwise comparisons rather than scalar rewards. Each round uses two independent Thompson draws to concentrate queries on informative duels and guides the search toward a *Copeland-optimal* prompt.

Notations. For each pair (p_i, p_j) , let W_{ij} denote the current number of wins of p_i over p_j , and let $N_{ij} = W_{ij} + W_{ji}$ be the total number of duels. The Bayesian posterior for the probability that p_i beats p_j is modeled as $\theta_{ij} \sim \text{Beta}(W_{ij} + 1, W_{ji} + 1)$. For $\alpha > 0$, the corresponding upper and lower confidence bounds are defined as

$$u_{ij}(t) = \frac{W_{ij}}{N_{ij}} + \sqrt{\frac{\alpha \log t}{\max\{1, N_{ij}\}}},$$

$$l_{ij}(t) = \frac{W_{ij}}{N_{ij}} - \sqrt{\frac{\alpha \log t}{\max\{1, N_{ij}\}}}.$$

Per-round Procedure ($t = 1, 2, \dots, T$):

1. *First Prompt Selection.* Compute an optimistic Copeland score for each prompt:

$$\hat{\zeta}_i(t) = \frac{1}{K-1} \sum_{j \neq i} \mathbf{1}\{u_{ij}(t) \geq 0.5\}$$

Keep only prompts with the maximum score in a set $\zeta(t) = \{i \mid \hat{\zeta}_i(t) = \max_k \hat{\zeta}_k(t)\}$. For each $i \in \zeta(t)$, draw independent samples

$\theta_{ij}^{(1)}$ and count

$$s_i = \sum_{j \neq i} \mathbf{1}\{\theta_{ij}^{(1)} \geq 0.5\}$$

Select prompt $i^* = \arg \max_{i \in \zeta(t)} s_i$.

2. *Second Prompt Selection.* Restrict to only uncertain opponents for i^* ,

$$S_{i^*}(t) = \{j \neq i^* : l_{i^*j}(t) \leq 0.5\}$$

Draw independent samples $\theta_{ji^*}^{(2)}$ and select

$$j^* = \arg \max_{j \in S_{i^*}(t)} \theta_{ji^*}^{(2)}$$

3. *Duel and update.* Judge prompts (p_{i^*}, p_{j^*}) , record the winner, and update $W_{i^*j^*}$.

3.2 Efficient Prompt Discovery via Top-Performer Guided Mutation

While D-TS identifies the Copeland-optimal prompt within a fixed pool, the broader goal of PDO is to locate the global optimum p^* over a combinatorially large prompt space. Because exhaustive search is infeasible, PDO, inspired by evolutionary prompt optimization techniques (Fernando et al., 2023; Guo et al., 2023), incrementally expands the candidate pool by mutating top-performing prompts and introducing mutation tips that promote variation and diversity without straying too far from high-performing prompts.

Mutation Procedure. At round t , let the current pool be $\mathcal{P}_t = \{p_1, \dots, p_{K_t}\}$ with the empirical Copeland scores $\hat{\mathbf{C}}_t := (\hat{C}_t(p))_{p \in \mathcal{P}_t}$. The procedure is:

1. *Selection:* Choose the top-performing prompt $p_t^* = \arg \max_{p \in \mathcal{P}_t} \hat{C}_t$.
2. *Mutation:* Generate a variant p_{new} of p_t^* via template edits, text-gradient guided changes, or LLM-assisted rewrites.
3. *Expansion:* Update pool $\mathcal{P}_{t+1} = \mathcal{P}_t \cup \{p_{\text{new}}\}$.

3.3 Theoretical insights of PDO

PDO can be viewed as a two-level optimization scheme that alternates *selection* and *discovery*. Within any fixed candidate pool, D-TS provides a principled way to allocate pairwise comparisons and is known to achieve $O(K^2 \log T)$ expected Copeland regret in general settings, implying vanishing average regret and convergence to the Copeland-winner set as the comparison budget grows (Wu and Liu, 2016).

To move beyond a fixed pool, PDO periodically expands the candidate set by mutating the current

Algorithm 1: Prompt Duel Optimizer

Require : \mathcal{P}_0 (initial pool); judge \mathcal{J} ; batch size m ; total rounds T ; mutation period M

- 1 $\mathcal{P} \leftarrow \mathcal{P}_0$; initialize $W, N \in \mathbb{R}^{|\mathcal{P}| \times |\mathcal{P}|}$ to zeros;
- 2 **for** $t \leftarrow 1$ **to** T **do**
- 3 $(p_i, p_j) \leftarrow \text{D-TS}(\mathcal{P}, W, N)$;
 // prompt selection via D-TS
- 4 $(w_i, w_j) \leftarrow \text{DUEL}(p_i, p_j, m, \mathcal{J})$;
 // get win counts
- 5 $W[i, j] \leftarrow W[i, j] + w_i$;
- 6 $W[j, i] \leftarrow W[j, i] + w_j$;
- 7 $N[i, j] \leftarrow N[i, j] + m$;
- 8 $N[j, i] \leftarrow N[j, i] + m$;
- 9 **if** $t \bmod M = 0$ **then**
- 10 $p_t^* \leftarrow \text{CURRENTBEST}(\mathcal{P}, W, N)$;
 // Copeland rank
- 11 $p_{\text{new}} \leftarrow \text{MUTATE}(p_t^*)$;
 // generate new prompts via
 top-performer guided
 mutation
- 12 $\mathcal{P} \leftarrow \mathcal{P} \cup \{p_{\text{new}}\}$;
- 13 Expand W, N with zero row/column
 for p_{new} ;
- 14 **return** $\text{CURRENTBEST}(\mathcal{P}, W, N)$;
 // final Copeland winner

Copeland leader. Under a standard smoothness assumption that prompt utility varies gradually under local edits (e.g., L -Lipschitz with respect to a prompt distance), this strategy concentrates new candidates in increasingly competitive neighborhoods, analogous to the “zooming-in” intuition in Lipschitz/metric bandits (Kleinberg et al., 2008).

Taken together, PDO separates prompt optimization into (i) sample-efficient identification of the best prompt among current candidates (via D-TS) and (ii) targeted exploration that increases coverage of near-optimal regions (via top-performer mutation), which helps explain its empirical efficiency under limited judge budgets.

4 Experiment and Results

4.1 Experimental Setup

Datasets. In this section, we conduct experiments to evaluate PDO on both closed-ended multiple-choice tasks and open-ended QA tasks. For the multiple-choice setting, we select 16 tasks from Big-Bench-Hard (Suzgun et al., 2022), using

accuracy as the evaluation metric. For the open-ended QA setting, we consider four task categories from MS-MARCO (Bajaj et al., 2016), where the final evaluation metric is an integer score between 1 and 5, assigned by an LLM judge comparing the model’s output with the ground-truth answer provided by the original dataset.

Across both task types, for the results reported in this section, we randomly split the data into development and test sets with a 50/50 split ratio. We report the test set performance averaged over 10 runs. The *Llama-3.3-70B-Instruct* model is used for prompt generation, preference judging, and final evaluation. Detailed descriptions of the datasets and the experimental setup of PDO are provided in Appendix B.

LLM Judge Design. For multiple-choice tasks, we require the LLM to produce both an answer and the accompanying reasoning given the instruction prompt. We then apply a dual-judge approach: if two prompts yield different answers, the LLM judge selects the prompt with the correct answer; if they yield the same answer, the judge decides based on the quality of reasoning. For open-ended tasks, the LLM judge design is more straightforward: each pair of responses is evaluated on accuracy, completeness, relevance, and clarity, consistent with the criteria of the final evaluation metric. To mitigate position bias, the order of the two outputs in each prompt pair is randomized before being fed into the judge template. A detailed rationale and analysis of the LLM preference judge design are provided in Section 5 and Appendix D.

Baselines. PDO is designed for prompt optimization without access to ground-truth labels or external references. A directly comparable baseline is **SPO** (Xiang et al., 2025), which similarly uses an LLM judge to iteratively compare the outputs of two prompts and select the winning prompt. We also include classical prompting techniques including chain-of-thought **COT** (Wei et al., 2022) and plan-and-solve **PoS** (Wang et al., 2023).

For the Big-Bench-Hard (BBH) dataset, we also evaluate PDO under a supervised setting to enable comparison with popular supervised APO methods, including **APE** (Zhou et al., 2022), **OPRO** (Yang et al., 2024), and **Breeder** (Fernando et al., 2023). The optimization itself remains label-free; labels are introduced only in the final stage to select the prompt that achieves the highest development-set accuracy among candidates generated through ini-

Method	Causal	Date	DisambigQA	Formal	Geometric	Hyperbaton	Logical-5	Logical-7
No prompt	<u>0.661</u> ± 0.044	0.854 ± 0.024	0.698 ± 0.047	<u>0.739</u> ± 0.031	0.434 ± 0.036	0.900 ± 0.020	0.785 ± 0.018	<u>0.739</u> ± 0.033
CoT	0.653 ± 0.042	0.877 ± 0.019	0.720 ± 0.039	0.725 ± 0.027	0.422 ± 0.028	0.891 ± 0.023	0.761 ± 0.027	0.726 ± 0.025
PoS	0.652 ± 0.037	0.878 ± 0.019	0.698 ± 0.043	<u>0.739</u> ± 0.027	0.403 ± 0.030	0.896 ± 0.024	<u>0.798</u> ± 0.032	0.750 ± 0.026
SPO	0.655 ± 0.033	<u>0.884</u> ± 0.017	<u>0.725</u> ± 0.057	0.738 ± 0.018	0.650 ± 0.069	0.886 ± 0.031	0.787 ± 0.031	0.721 ± 0.026
PDO (ours)	0.681 ± 0.040	0.918 ± 0.014	0.738 ± 0.050	0.744 ± 0.030	<u>0.598</u> ± 0.073	0.910 ± 0.029	0.804 ± 0.034	0.711 ± 0.019
Method	Navigate	Penguins	Salient	Snarks	Tracking-5	Tracking-7	Tracking-3	Web of Lies
No prompt	0.869 ± 0.013	0.915 ± 0.018	<u>0.698</u> ± 0.020	0.823 ± 0.027	0.695 ± 0.033	0.499 ± 0.020	0.890 ± 0.019	0.766 ± 0.020
CoT	<u>0.878</u> ± 0.016	0.915 ± 0.023	0.709 ± 0.021	<u>0.833</u> ± 0.023	0.724 ± 0.046	0.532 ± 0.025	<u>0.904</u> ± 0.019	0.796 ± 0.022
PoS	0.866 ± 0.019	0.910 ± 0.027	0.693 ± 0.025	0.816 ± 0.026	<u>0.725</u> ± 0.030	0.538 ± 0.034	0.888 ± 0.019	<u>0.861</u> ± 0.019
SPO	0.874 ± 0.035	<u>0.934</u> ± 0.025	0.662 ± 0.038	0.820 ± 0.046	0.692 ± 0.046	<u>0.543</u> ± 0.026	0.826 ± 0.087	0.818 ± 0.043
PDO (ours)	0.900 ± 0.023	0.937 ± 0.034	0.681 ± 0.032	0.840 ± 0.039	0.796 ± 0.084	0.641 ± 0.089	0.930 ± 0.046	0.942 ± 0.040

Table 1: Test results on 16 BBH tasks averaged over 10 runs. For PDO, we report the test accuracy of the prompt selected by the highest Copeland score. PDO is compared with baselines that assume **no access** to ground-truth labels. The best performance is shown in bold, and the second-best is underlined.

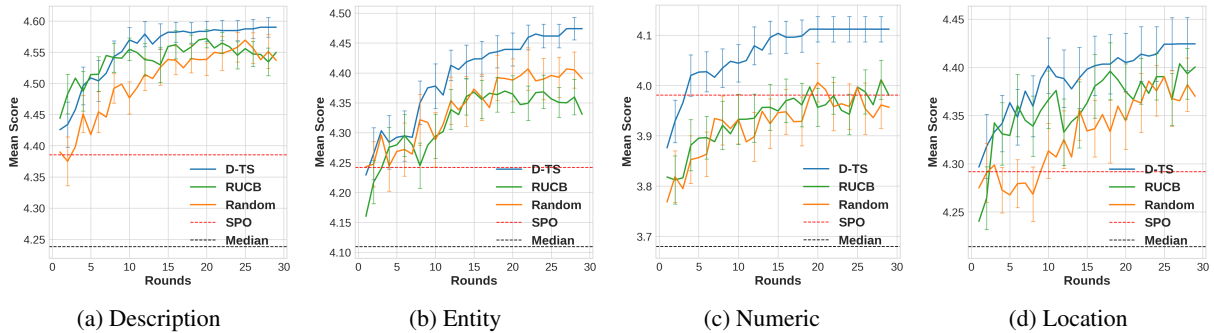


Figure 2: Test performance of the winning prompt on the four MS-MARCO tasks. Each curve shows the mean score of the current Copeland leader over rounds, with 50 duels per round. PDO with D-TS consistently outperforms RUCB and Random sampling, and surpasses the SPO baseline within a few rounds. Gray lines indicate the median test score across all prompts generated by PDO.

tial proposals and prompt mutations. Results for the without-label and with-label settings are reported in Tables 1 and 6, respectively.

4.2 Benchmark Results

Multiple-choice Tasks Performance. As shown in Table 1, using only preference signals from the LLM judge and selecting the winner prompt by Copeland scores, PDO achieves the highest evaluation accuracy on **13/16** tasks. Notable gains include *Tracking-7* (0.641 vs. 0.543, **+9.8pp**) and *Web of Lies* (0.942 vs. 0.861, **+8.1pp**).

Open-ended QA Tasks. We evaluate four MS-MARCO tasks (*Description*, *Entity*, *Numeric*, and *Location*), starting from a pool of $|\mathcal{P}| = 50$ instructions. At each round t , we snapshot the win matrix W_t , compute Copeland scores, and report the test performance of the current Copeland winner. Figure 2 reports averages over 30 independent runs, comparing D-TS with the dueling-bandit alternative Relative Upper Confidence Bound (RUCB; Zoghi et al. 2013), uniform Random sampling, and the SPO baseline. Across all tasks, D-TS achieves

the highest scores and converges more quickly than RUCB and Random. The horizontal reference lines mark the median score of all prompts generated by PDO (gray) and the SPO baseline (red). Within just a few rounds, D-TS surpasses both reference lines and maintains this advantage over RUCB and Random. These results highlight the *sample efficiency* of D-TS: it consistently identifies stronger prompts faster and more reliably than random sampling or alternative dueling-bandit methods.

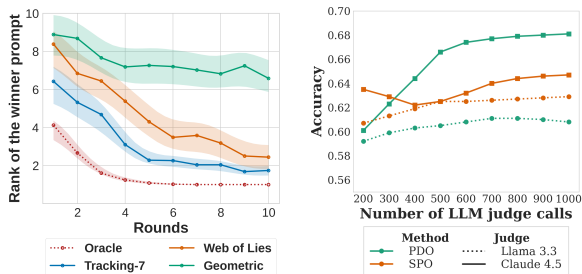
5 LLM Judge Analysis

In the experiments reported above, we use two types of LLM judges. The first is the *preference judge* used during LLM prompt optimization to compare outputs from two candidate prompts. The second is the *evaluation judge*, which provides the final score for the open-ended MS-MARCO task. In this section, we present additional experiments to analyze the noise and potential bias introduced by both judge types.

5.1 Preference Judge

Correlation Between Judge Noise and Task Performance. We proxy preference judge effectiveness for the 16 BBH tasks by the test accuracy gap between Table 1 (prompt chosen by Copeland score, label-free) and Table 6 (prompt chosen by development-set accuracy). In our results, *Tracking-7* and *Web of Lies* show small gaps and achieve top performance under both selection criteria, whereas *Geometric* exhibits a large discrepancy of **11.4** percentage points.

To investigate further, for each task we fix $|\mathcal{P}| = 20$ instruction prompts with varying accuracy, run D-TS with the LLM judge, and compare it against an *oracle* judge that always selects the higher-accuracy prompt in each duel. At round t , we record the ground-truth accuracy rank of the current Copeland leader. Figure 3a shows that the oracle converges to the best prompt by round 4; with the LLM judge, *Tracking-7* steadily improves to rank 2 and *Web of Lies* approaches rank ≈ 2.5 , while *Geometric* remains around ranks 6–8 across rounds. These trends confirm that judge reliability is closely related to the performance gaps observed in Tables 1 and 6.



(a) Judge noise across different BBH tasks. (b) Noise reduction on *Geometric* using Claude 4.5.

Figure 3: (a) The accuracy of the LLM judge varies across tasks, with some tasks (e.g., *Geometric*) showing persistent misjudgments compared to an oracle judge. (b) Using a frontier LLM judge (Claude 4.5) helps reduce the impact of noisy judgements.

Reducing Judge Noise using frontier model.

We hypothesize that a key source of judge noise on BBH tasks arises from heterogeneous task difficulty: when two prompts produce different answers, the preference judge must determine which output is correct, and this decision becomes unreliable on harder tasks for judges with limited reasoning ability. A natural remedy is to replace the current Llama-3.3 preference judge with a stronger

frontier model that better approximates an oracle, therefore reducing decision noise.

Table 2 reports the selection accuracy of the preference judge on examples where the two candidate answers differ. With Llama-3.3-70B as the judge, accuracy is high on *Tracking-7* and *Web of Lies* but only marginally above chance on *Geometric*, consistent with the trend in Figure 3a. Substituting Claude-4.5-Sonnet as the preference judge substantially improves selection accuracy on *Geometric* and yields near-oracle performance on the other two tasks. As a result, when replacing the preference judge with Claude 4.5, the test accuracy of the prompts on *Geometric* found by PDO exceeds that of SPO across LLM judge API budgets as reported in Figure 3b. We report additional results on three other BBH tasks in Figure 7.

Judge Model	Geometric	Tracking-7	Web of Lies
Llama-3.3-70B	0.59	0.89	0.85
Claude-4.5-Sonnet	0.87	0.98	0.99

Table 2: Preference-judge selection accuracy on BBH instances where two candidate prompts produce different answers (1 = always selects the correct answer; 0.5 = random guessing).

5.2 Evaluation Judge

Cross-family Evaluation. For the MS MARCO results in Section 4, we use the same Llama 3.3 model both to guide prompt optimization and to score the final prompt, which may introduce judge circularity (i.e., the optimized prompt is favored by the same model used for evaluation). To probe this issue beyond a single prompt pair, we form two prompt sets by sampling uniformly at random: 10 prompts produced by PDO and 10 prompts produced by SPO (all with their original Llama 3.3 judge scores recorded during optimization). To avoid additional inference variability, we use cached model outputs for each prompt on the same evaluation set and re-score these fixed outputs against the reference answers using judges from different model families.

Table 3 reports cross-family re-scoring results across the 10 sampled prompts per method. PDO remains higher than SPO under every judge, with consistently positive margins, suggesting that the observed improvement is not merely an artifact of circularity with the original Llama 3.3 judge.

In Appendix A, we include a small-scale human agreement study testing whether the LLM judge

evaluation on MS-MARCO aligns with human preferences on examples where PDO and SPO outputs receive different scores.

Judge Model	PDO	SPO	Δ
Llama 3.3 (original)	4.68 \pm 0.06	4.50 \pm 0.07	0.18
Mistral-Large	4.61 \pm 0.07	4.49 \pm 0.07	0.12
GPT-4o	4.60 \pm 0.06	4.41 \pm 0.06	0.19
Claude 3.5 Sonnet	4.45 \pm 0.08	4.25 \pm 0.08	0.20
Claude 4.5 Sonnet	4.58 \pm 0.06	4.49 \pm 0.06	0.09

Table 3: Cross-family LLM-judge re-evaluation on MS MARCO (Description).

6 Ablation Study

Cost Analysis. We make two kinds of LLM calls in prompt optimization: inference calls to generate a prompt’s outputs and judge calls to compare outputs pairwise. Because our evaluation is empirical, we assess sample efficiency by comparing methods under the same judge-call budget.

Figure 2 shows that D-TS consistently finds better prompts than random sampling given the same number of judge calls. We also report mean wall-clock time on MS MARCO for SPO with its default 20 rounds and for PDO early-stopped at 10 rounds (Table 4). We use 10 rounds for PDO because Figure 2 indicates that by round 10, the prompt selected by PDO already surpasses SPO in test accuracy.

SPO is faster per run largely because it compares only two prompts per round and immediately eliminates the loser. However, without an explicit exploration–exploitation strategy, this procedure can be brittle under noisy comparisons, which helps explain its weaker performance.

Practically, as discussed in Section 5.1, we control judge noise by using a strong but expensive model (e.g., Claude-4.5) to guide the search, and then deploying the optimized prompt on a lower-cost model (e.g., Llama-3.3) that serves production traffic. In this setting, D-TS is especially effective because it achieves better prompts with fewer expensive judge calls.

Method	Description	Entity	Numeric	Location
PDO	4.39	3.39	3.83	3.10
SPO	6.78	6.74	6.17	7.45

Table 4: Mean wall-clock time (minutes) on MS MARCO across four prompt tasks.

Prompt Mutation. PDO incrementally expands the prompt pool by mutating top-performing




Figure 4: Effect of top-performer–guided mutation on *Web of Lies* (left) and *Tracking-7* (right).

prompts. We visualize its effect on *Web of Lies* and *Tracking-7*. In the *mutate* variant, at rounds 10 and 20 we prune the 10 lowest–Copeland-score prompts and replace them with 10 new candidates generated by mutating the current Copeland leader. The *non_mutate* baseline keeps the original 20 prompts fixed throughout. At each round, we report the ground-truth accuracy of the current Copeland leader. As shown in Figure 4, both methods behave similarly before round 10, but mutation yields consistently higher accuracy afterward, surpassing the limits of the fixed pool and navigating toward better prompt regions.

We additionally report results across all 16 BBH tasks in Table 7. In total, 8 out of 16 tasks show statistically significant gains under mutation, and mutation never underperforms relative to the non-mutation baseline.

Pairwise Preference vs. Pointwise Scoring.

PDO relies on the *pairwise preferences* of an LLM judge to select the winning prompt. To highlight the benefit of this formulation, we compare it against a *pointwise scoring* judge, where each prompt is evaluated independently on the development set by assigning a numeric score (1–5) without access to ground-truth answers. The prompt with the highest average score is then selected. We evaluate both strategies on the same fixed pool of 50 candidates from MS-MARCO in 4.2 and report the test mean score of the selected prompt. Figure 5 shows that preference judgement consistently outperforms pointwise scoring in 7 of 8 model–task combinations across two judge models. Overall, these results are consistent with the LLM-as-judge literature: pairwise comparisons yield more stable decisions than direct pointwise scoring by reducing dependence on calibration (Liu et al., 2024).

Round	Method	Description	Entity	Numeric	Location
5	D-TS	4.52	<u>4.38</u>	4.00	4.32
	Self-Sparring	<u>4.50</u>	4.40	4.00	<u>4.30</u>
	RMED	4.47	4.37	<u>3.94</u>	4.29
	RUCB	4.46	4.32	3.92	4.26
	MergeRUCB	4.44	4.34	3.88	4.28
	Random	4.40	4.36	3.89	4.20
10	D-TS	4.57	<u>4.44</u>	4.07	4.41
	Self-Sparring	<u>4.56</u>	4.46	<u>4.06</u>	<u>4.40</u>
	RMED	4.53	4.42	4.02	4.39
	RUCB	4.52	4.38	4.00	4.36
	MergeRUCB	4.50	4.40	3.95	4.38
	Random	4.45	4.41	3.97	4.29
15	D-TS	4.61	<u>4.48</u>	4.11	4.45
	Self-Sparring	<u>4.59</u>	4.50	4.11	<u>4.43</u>
	RMED	4.57	4.45	<u>4.05</u>	4.42
	RUCB	4.55	4.40	4.03	4.39
	MergeRUCB	4.53	4.42	3.98	4.41
	Random	4.48	4.44	4.00	4.33

Table 5: Comparison of exploration strategies on MS-MARCO across optimization rounds in PDO. We report the mean test score of the current Copeland leader over rounds, selected by different exploration strategies. For each {round, dataset}, the best score is shown in **bold** and the second best is underlined.




Figure 5: Comparison of pairwise preference used in PDO and pointwise scoring for prompt selection on MS-MARCO. Results indicate that pairwise preference leads to higher-performing prompts across model sizes (7 out of 8 cases).

Analysis of Exploration Strategies The dueling-bandit literature spans several algorithmic families. A useful high-level division is between *non-stochastic* index policies, which select the next duel deterministically based on confidence or divergence indices, and *stochastic* policies, which randomize comparisons via posterior sampling. To justify the use of D-TS as the exploration strategy in PDO, we conduct additional experiments comparing two UCB-style index methods (RUCB (Zoghi et al., 2013) and MergeRUCB (Zoghi et al., 2015)), a KL-divergence index method (RMED (Komiya et al., 2015)), and two posterior-sampling methods (D-TS (Wu and Liu, 2016) and Self-Sparring (Sui et al., 2017)).

Table 5 shows monotonic improvement from Round 5 to 15 for all methods. Under the same judge budget, posterior-sampling approaches are consistently strongest: D-TS achieves the best performance in most (round, task) cells, while Self-Sparring remains highly competitive and is the top-performing method on ENTITY at every reported round. RMED consistently falls between the TS-based methods and the UCB-style baselines—outperforming RUCB and MergeRUCB but remaining slightly below posterior sampling—while random selection performs worst overall.

We conjecture that this performance gap reflects characteristics of the prompt-optimization setting: a short optimization horizon with noisy preference feedback, where many prompt pairs are difficult to reliably separate and the effective notion of “best”

aligns more closely with overall win frequency than with a clearly dominant Condorcet winner. In such regimes, posterior sampling can act as an *anytime*, uncertainty-aware heuristic that continues to explore multiple plausibly strong prompts without over-investing in resolving near-ties. By contrast, index-based methods (both UCB- and divergence-based) may allocate a larger fraction of a limited budget to shrinking confidence sets or satisfying information constraints, which can be less effective when the horizon is small and the judge signal is noisy.

7 Related Work

APO Methods: from Supervised to Label-free.

Traditional prompt optimization methods rely heavily on supervised signals from labeled validation sets, where the typical pipeline involves generating, mutating, and scoring prompts (Zhou et al., 2022; Fernando et al., 2023; Pryzant et al., 2023; Guo et al., 2023; Schulhoff et al., 2024). More recently, studies have reduced supervision requirements (Xiang et al., 2025; Chen et al., 2024; Madaan et al., 2023; Cheng et al., 2023; Dong et al., 2022). Notably, Self-Supervised Prompt Optimization (SPO) (Xiang et al., 2025) eliminates external references by using output-vs-output comparisons, iteratively generating, executing, and comparing prompts through pairwise LLM judgments. However, SPO follows a *greedy hill-climbing loop* without a principled exploration-exploitation strategy, limiting its ability to allocate comparisons efficiently and to robustly identify high-performing prompts.

Bandit-Based Prompt Optimization. The connection between prompt optimization and multi-armed bandits (MAB) has recently gained attention, as prompts can naturally be viewed as arms. OPTS (Ashizawa et al., 2025) formulates prompt strategy selection as a bandit problem with Thompson sampling. TRIPLE (Shi et al., 2024) casts prompt optimization as fixed-budget best-arm identification, assuming a predefined set of prompts with known scores. However, both approaches require labeled validation sets for scoring. APOHF (Lin et al., 2025) connects preference-based prompt optimization to dueling bandits, but assumes human-annotated pairwise preferences that are impractical at scale. PDO retains the dueling-bandit formulation while replacing human preferences with LLM-judge comparisons to align with realistic prompt optimization settings.

8 Conclusion

In this paper, we introduce PDO, a label-free prompt optimization method by framing the problem as a dueling bandit guided by LLM preference feedback. By combining Double Thompson Sampling (D-TS) with top-performer-guided mutation, PDO adaptively searches for informative prompt comparisons to identify high-performing prompts. Experiments on BBH and MS-MARCO datasets show that PDO consistently achieves competitive performance against baseline methods across tasks. We analyze judge noise and potential circularity, and demonstrate that PDO’s gains persist under alternative judge families. More broadly, our results suggest that LLM pairwise preference feedback can serve as a practical supervision signal for prompt optimization when labeled data are unavailable. Beyond the single-turn settings considered here, PDO provides a flexible foundation for extending label-free prompt optimization to more complex scenarios, such as multi-turn interactions and prompt sets for agentic systems.

Limitation

While our dueling-bandit framework provides a practical and sample-efficient approach to prompt optimization without ground-truth labels, it also has inherent limitations. The capability of the preference judge plays a critical role: depending on task complexity, the judge may be noisy. Stronger foundation models tend to yield more reliable preferences but incur substantially higher cost, whereas smaller or less specialized models may introduce greater noise, particularly on domain-specific tasks. Because the algorithm optimizes for the judge’s notion of quality rather than the true task metric, it may favor stylistic patterns that align with the judge’s preferences.

To mitigate these concerns, we include a detailed discussion and ablation studies on LLM judge design, along with a small-scale human evaluation, in Section 5 and Appendix D. However, we were unable to conduct broader experiments across a wider range of tasks. Despite these challenges, our formulation provides a principled and scalable starting point for label-free prompt optimization using preference feedback, and we hope it motivates future work on improving the alignment between LLM judges and task objectives.

References

- Shipra Agrawal and Navin Goyal. 2012. [Analysis of thompson sampling for the multi-armed bandit problem](#). In *Proceedings of the 25th Annual Conference on Learning Theory*, volume 23 of *Proceedings of Machine Learning Research*, pages 39.1–39.26, Edinburgh, Scotland. PMLR.
- Rin Ashizawa, Yoichi Hirose, Nozomu Yoshinari, Kento Uchida, and Shinichi Shirakawa. 2025. [Bandit-based prompt design strategy selection improves prompt optimizers](#). In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 20799–20817, Vienna, Austria. Association for Computational Linguistics.
- Payal Bajaj, Daniel Campos, Nick Craswell, Li Deng, Jianfeng Gao, Xiaodong Liu, Rangan Majumder, Andrew McNamara, Bhaskar Mitra, Tri Nguyen, Mir Rosenberg, Xia Song, Alina Stoica, Saurabh Tiwary, and Tong Wang. 2016. [Ms marco: A human generated machine reading comprehension dataset](#). *arXiv preprint arXiv:1611.09268*.
- Viktor Bengs, Róbert Busa-Fekete, Adil El Mésaoudi-Paul, and Eyke Hüllermeier. 2021. [Preference-based online learning with dueling bandits: A survey](#). *Journal of Machine Learning Research*, 22(7):1–108. Open access.
- Yongchao Chen, Jacob Arkin, Yilun Hao, Yang Zhang, Nicholas Roy, and Chuchu Fan. 2024. [PRompt optimization in multi-step tasks \(PROMST\): Integrating human feedback and heuristic-based sampling](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 3859–3920, Miami, Florida, USA. Association for Computational Linguistics.
- Jiale Cheng, Xiao Liu, Kehan Zheng, Pei Ke, Hongning Wang, Yuxiao Dong, Jie Tang, and Minlie Huang. 2023. [Black-box prompt optimization: Aligning large language models without model training](#). *ArXiv*, abs/2311.04155.
- Yihong Dong, Kangcheng Luo, Xue Jiang, Zhi Jin, and Ge Li. 2022. [Pace: Improving prompt with actor-critic editing for large language model](#). In *arXiv preprint arXiv:2212.XXXX*.
- Chrisantha Fernando, Dylan Banarse, Henryk Michalewski, Simon Osindero, and Tim Rocktäschel. 2023. [Promptbreeder: Self-referential self-improvement via prompt evolution](#). *arXiv preprint arXiv:2309.16797*.
- Jiawei Gu, Xuhui Jiang, Zhichao Shi, Hexiang Tan, Xuehao Zhai, Chengjin Xu, Wei Li, Yinghan Shen, Shengjie Ma, Honghao Liu, Saizhuo Wang, Kun Zhang, Yuanzhuo Wang, Wen Gao, Lionel M. Ni, and Jian Guo. 2024. [A survey on llm-as-a-judge](#). *arXiv preprint arXiv:2411.15594*.
- Qingyan Guo, Rui Wang, Junliang Guo, Bei Li, Kaitao Song, Xu Tan, Guoqing Liu, Jiang Bian, and Yujie Yang. 2023. [Evoprompt: Connecting llms with evolutionary algorithms yields powerful prompt optimizers](#). *arXiv preprint arXiv:2309.08532*.
- Horace He and Thinking Machines Lab. 2025. [Defeating nondeterminism in llm inference](#). *Thinking Machines Lab: Connectionism*. <https://thinkingmachines.ai/blog/defeating-nondeterminism-in-llm-inference/>.
- Robert Kleinberg, Aleksandrs Slivkins, and Eli Upfal. 2008. [Multi-armed bandits in metric spaces](#). In *Proceedings of the 40th Annual ACM Symposium on Theory of Computing (STOC)*, pages 681–690. ACM.
- Junpei Komiyama, Junya Honda, Hisashi Kashima, and Hiroshi Nakagawa. 2015. [Regret lower bound and optimal algorithm in dueling bandit problem](#). In *Proceedings of The 28th Conference on Learning Theory (COLT)*, pages 1141–1154.
- Xiaoqiang Lin, Zhongxiang Dai, Arun Verma, See-Kiong Ng, Patrick Jaillet, and Bryan Kian Hsiang Low. 2025. [Prompt optimization with human feedback](#).
- Yang Liu, Dan Iter, Yichong Xu, Shuhang Wang, Ruochen Xu, and Chenguang Zhu. 2023. [G-eval: NLG evaluation using gpt-4 with better human alignment](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 2511–2522, Singapore. Association for Computational Linguistics.
- Yinhong Liu, Han Zhou, Zhijiang Guo, Ehsan Shareghi, Ivan Vulić, Anna Korhonen, and Nigel Collier. 2024. [Aligning with human judgement: The role of pairwise preference in large language model evaluators](#). In *Proceedings of the First Conference on Language Modeling (COLM 2024)*, page –.
- Aman Madaan, Niket Tandon, Prakhar Gupta, Skyler Hallinan, Luyu Gao, Sarah Wiegrefe, Uri Alon, Nouha Dziri, Shrimai Prabhumoye, Yiming Yang, Sean Welleck, Bodhisattwa Prasad Majumder, Shashank Gupta, Amir Yazdanbakhsh, and Peter Clark. 2023. [Self-refine: Iterative refinement with self-feedback](#). *ArXiv*, abs/2303.17651.
- Krista Opsahl-Ong, Michael J. Ryan, Josh Purtell, David Broman, Christopher Potts, Matei Zaharia, and Omar Khattab. 2024. [Optimizing instructions and demonstrations for multi-stage language model programs](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 9340–9366, Miami, Florida, USA. Association for Computational Linguistics.
- Konstantina Palla, José Luis Redondo García, Claudia Hauff, Francesco Fabbri, Henrik Lindström, Daniel R. Taber, Andreas Damianou, and Mounia Lalmas. 2025. [Policy-as-prompt: Rethinking content moderation in the age of large language models](#). In *Proceedings of the ACM Conference on Fairness, Accountability, and Transparency (FAccT 2025)*, pages 840–854.

- Reid Pryzant, Dan Iter, Jerry Li, Yin Lee, Chenguang Zhu, and Michael Zeng. 2023. [Automatic prompt optimization with “gradient descent” and beam search](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 7957–7968. Association for Computational Linguistics.
- Kiran Ramnath, Kang Zhou, Sheng Guan, Soumya Smruti Mishra, Xuan Qi, Zhengyuan Shen, Shuai Wang, Sangmin Woo, Sullam Jeoung, Yawei Wang, Haozhu Wang, Han Ding, Yuzhe Lu, Zhichao Xu, Yun Zhou, Balasubramaniam Srinivasan, Qiaojing Yan, Yueyan Chen, Haibo Ding, and 2 others. 2025. [A systematic survey of automatic prompt optimization techniques](#). *arXiv preprint arXiv:2502.16923v2*.
- Alexander Ratner, Stephen H. Bach, Henry Ehrenberg, Jason Fries, Sen Wu, and Christopher Ré. 2017. [Snorkel: Rapid training data creation with weak supervision](#). *Proceedings of the VLDB Endowment*, 11(3):269–282.
- Samuel Schulhoff and 1 others. 2024. [A systematic survey of prompting techniques](#). *arXiv preprint arXiv:2406.06608*.
- Chengshuai Shi, Kun Yang, Jing Yang, and Cong Shen. 2024. [Efficient prompt optimization through the lens of best arm identification](#). In *Neural Information Processing Systems*.
- Yanan Sui, Vincent Zhuang, Joel W. Burdick, and Yisong Yue. 2017. Multi-dueling bandits with dependent arms. In *Proceedings of the 33rd Conference on Uncertainty in Artificial Intelligence (UAI)*.
- Jiuding Sun, Chantal Shaib, and Byron C. Wallace. 2023. [Evaluating the zero-shot robustness of instruction-tuned language models](#). *arXiv preprint arXiv:2306.11270*.
- Mirac Suzgun, Nathan Scales, Nathanael Schärli, Sebastian Gehrmann, Yi Tay, Hyung Won Chung, Aakanksha Chowdhery, Quoc V Le, Ed H Chi, Denny Zhou, , and Jason Wei. 2022. [Challenging big-bench tasks and whether chain-of-thought can solve them](#). *arXiv preprint arXiv:2210.09261*.
- Patrick Wagner. 2024. [Using llms for text classification](#). *Medium*.
- Lei Wang, Wanyu Xu, Yihuai Lan, Zhiqiang Hu, Yunshi Lan, Roy Ka-Wei Lee, and Ee-Peng Lim. 2023. [Plan-and-solve prompting: Improving zero-shot chain-of-thought reasoning by large language models](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2609–2634, Toronto, Canada. Association for Computational Linguistics.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed H. Chi, Quoc V. Le, and Denny Zhou. 2022. [Chain-of-thought prompting elicits reasoning in large language models](#). In *NeurIPS*.
- Huasen Wu and Xin Liu. 2016. [Double thompson sampling for dueling bandits](#). *Advances in neural information processing systems*, 29.
- Jinyu Xiang, Jiayi Zhang, Zhaoyang Yu, Xinbing Liang, Fengwei Teng, Jinhao Tu, Fashen Ren, Xiangru Tang, Sirui Hong, Chenglin Wu, and Yuyu Luo. 2025. [Self-supervised prompt optimization](#). *arXiv preprint arXiv:2502.06855*. Version v3, revised 21 August 2025.
- Chengrun Yang, Xuezhi Wang, Yifeng Lu, Hanxiao Liu, Quoc V. Le, Denny Zhou, and Xinyun Chen. 2024. [Large language models as optimizers](#). In *ICLR*.
- Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric P. Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica. 2023. [Judging llm-as-a-judge with mt-bench and chatbot arena](#). *arXiv preprint arXiv:2306.05685*.
- 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](#). *arXiv preprint arXiv:2211.01910*.
- Masrour Zoghi, Shimon Whiteson, and Maarten de Rijke. 2015. [Mergerucb: A method for large-scale online ranker evaluation](#). In *Proceedings of the Eighth ACM International Conference on Web Search and Data Mining (WSDM)*, pages 17–26. ACM.
- Masrour Zoghi, Shimon Whiteson, Rémi Munos, and Maarten de Rijke. 2013. [Relative upper confidence bound for the k-armed dueling bandit problem](#). *arXiv preprint arXiv:1312.3393*. Version 2, submitted December 2013.

A Small-scale Human Agreement Check

In addition to the cross-judge evaluation in Table 3, we examine whether the original Llama 3.3 MS-MARCO judge aligns with human preferences. We select 20 evaluation examples (5 from each of the four MS-MARCO tasks) for which the cached PDO and SPO outputs receive different Llama 3.3 scores under the 1–5 rubric (e.g., 5 vs. 4). We then recruit 10 non-author volunteer raters (graduate-student participants from a university population) to perform a blinded pairwise comparison: for each example, raters are shown the input and two anonymized outputs (PDO vs. SPO), with left/right order randomized and prompt identities hidden, and are asked to choose the output they prefer under the same criteria as the Llama judge (accuracy, completeness, relevance, and clarity).

For each example, we compute an alignment rate as the fraction of raters whose preferred output matches the direction implied by the Llama scores, and we summarize the distribution of alignment rates across the 20 examples in Figure 6. Overall, this distribution suggests that the Llama-based evaluation signal is broadly consistent with human preferences on the subset of cases most relevant to distinguishing PDO from SPO, with a mean alignment ratio of 0.87.




Figure 6: Histogram of per-example human–Llama alignment rates on 20 MS MARCO examples, where PDO and SPO outputs receive different Llama 3.3 evaluation scores.

B Experiment Details

B.1 Hyperparameters

For the main experiments in Section 4.2, we initialize with 20 candidate prompts for BBH and 50 candidate prompts for MS-MARCO. Each experiment runs for 30 rounds, with 25 duels per round. At rounds 10 and 20, we apply prompt mutation

by selecting the top-3 prompts ranked by Copeland scores and generating 10 new prompts. At the same time, we prune the 10 lowest-ranked prompts by Copeland scores. We always select the prompt with the highest Copeland score as the winner, using the average win rate as a tiebreaker when multiple prompts share the same score. The D-TS parameter is fixed at $\alpha = 1.2$ throughout all experiments.

B.2 Prompt generation strategies

To generate the initial candidate prompts, we follow the approach in MiPROV2 (Opsahl-Ong et al., 2024), which first uses an LLM to construct a dataset summary from demonstration examples. The prompt generation process then incorporates this dataset summary, additional unlabelled demonstration examples, and randomly selected prompt tips from a predefined set to encourage diversity. In the mutation stage, we create new variants of the top-performing prompt by applying prompt tips to guide local edits, such as changing the tone, adjusting the wording, or appending synthetic few-shot examples.

B.3 Baselines

Unsupervised (label-free) methods SPO follows an iterative optimize–execute–evaluate loop. In each round, the method first samples n_{examples} fresh examples and evaluates the current prompt by running one answer-generation LLM call per example. A revised prompt is then proposed using a single LLM call conditioned on the observed cases. Both the current and revised prompts are evaluated on the same set of examples via an additional n_{examples} answer-generation calls. The two sets of outputs are subsequently compared using $n_{\text{judge_trials}}$ pairwise LLM judge calls, where each judge observes both outputs with randomized order to mitigate positional bias and votes for the preferred prompt. The revised prompt is accepted if it receives a majority of the judge votes; otherwise, the current prompt is retained. Overall, each round incurs $2n_{\text{examples}} + n_{\text{judge_trials}} + 1$ LLM calls.

With the default setting $n_{\text{examples}} = 3$ and $n_{\text{judge_trials}} = 3$, this corresponds to 10 calls per round, and we run 20 rounds to obtain the main results reported in Section 4.

Supervised methods For supervised APO methods, the typical pipeline involves generating, mutating, and scoring prompts. APE frames prompt design as black-box optimization, generating can-

didate instructions from input–output demonstrations and refining them through iterative search. **OPRO** treats the LLM itself as the optimizer: a meta-prompt with task descriptions and prior results guides the model to propose and evaluate new candidates in an iterative loop. **Breeder** applies an evolutionary approach, jointly evolving task-prompts and mutation-prompts through LLM-driven mutation and selection.

B.4 Datasets

BIG-bench Hard (BBH). BBH (Suzgun et al., 2022) is a curated subset of BIG-bench consisting of 23 reasoning-intensive tasks. It is commonly used to stress multi-step reasoning and symbolic manipulation. In our experiments, we evaluate on 16 BBH multiple-choice tasks, where LLaMA-3.3-70B shows non-trivial sensitivity to instruction prompts.

MS MARCO. MS MARCO (Bajaj et al., 2016) is a large-scale question answering dataset built from real Bing search queries, paired with human-written answers and linked passages. It supports QA, passage ranking, and related IR/NLP tasks. In our setting, we focus on four task categories—Description, Entity, Numeric, and Location—and adopt a 1–5 integer scoring scheme from an LLM judge that compares model outputs against the dataset’s ground-truth answers.

C Discussion on Evaluation Scope

We evaluate PDO on BBH and MS MARCO, which capture two practically important prompt-optimization regimes. BBH covers a broad set of real-world, prompt-based instruction-following *classification* tasks in a cold-start setting, where labeled data are scarce and costly to obtain at the outset. This setting is increasingly relevant in industry workflows such as content moderation (Palla et al., 2025), where many systems are transitioning from fully supervised ML/DL pipelines to prompt-based solutions whose performance is largely determined by the underlying policy prompt.

In contrast, MS MARCO represents a single-turn instruction/QA setting, mirroring scenarios where practitioners optimize a system prompt for a locally deployed LLM application. In such cases, tasks are often open-ended and evaluation criteria can be subjective, making an LLM judge a natural and scalable choice for preference-based comparisons.

We admit that extending PDO to multi-turn interactions, as well as jointly optimizing *sets* of prompts (e.g., for multi-agent systems), is a valuable and important direction for future work.

Method	Causal	Date	DisambigQA	Formal	Geometric	Hyperbaton	Logical-5	Logical-7
APE	0.680 \pm 0.044	0.892 \pm 0.019	0.730 \pm 0.043	0.747 \pm 0.022	0.670 \pm 0.072	0.940 \pm 0.013	0.822 \pm 0.041	0.721 \pm 0.035
OPRO	0.682 \pm 0.044	0.910 \pm 0.022	0.734 \pm 0.038	0.728 \pm 0.026	0.569 \pm 0.048	0.932 \pm 0.021	0.774 \pm 0.030	0.718 \pm 0.031
Breeder	0.683 \pm 0.027	0.898 \pm 0.026	0.745 \pm 0.038	0.746 \pm 0.027	0.684 \pm 0.048	0.932 \pm 0.021	0.778 \pm 0.022	0.724 \pm 0.018
PDO (ours)	0.680 \pm 0.033	0.915 \pm 0.024	0.740 \pm 0.040	0.754 \pm 0.022	0.712 \pm 0.058	0.948 \pm 0.017	0.809 \pm 0.028	0.733 \pm 0.024
Method	Navigate	Penguins	Salient	Snarks	Tracking-5	Tracking-7	Tracking-3	Web of Lies
APE	0.899 \pm 0.024	0.925 \pm 0.025	0.686 \pm 0.033	0.875 \pm 0.033	0.803 \pm 0.041	0.600 \pm 0.050	0.911 \pm 0.041	0.948 \pm 0.020
OPRO	0.882 \pm 0.032	0.919 \pm 0.025	0.687 \pm 0.023	0.888 \pm 0.021	0.833 \pm 0.083	0.662 \pm 0.073	0.924 \pm 0.035	0.938 \pm 0.045
Breeder	0.908 \pm 0.016	0.927 \pm 0.018	0.698 \pm 0.024	0.903 \pm 0.019	0.859 \pm 0.049	0.600 \pm 0.042	0.927 \pm 0.028	0.963 \pm 0.019
PDO (ours)	0.890 \pm 0.010	0.933 \pm 0.027	0.695 \pm 0.022	0.874 \pm 0.036	0.847 \pm 0.062	0.662 \pm 0.084	0.946 \pm 0.015	0.979 \pm 0.024

Table 6: We report test results for PDO when selecting the prompt with the highest development-set accuracy from the same candidate pools as in Table 1. This setting serves as a proxy for replacing the LLM preference judge with an oracle. We compare these results against supervised APO methods and find that PDO remains competitive with state-of-the-art prompt optimization baselines, achieving the best performance on 9 out of 16 tasks.

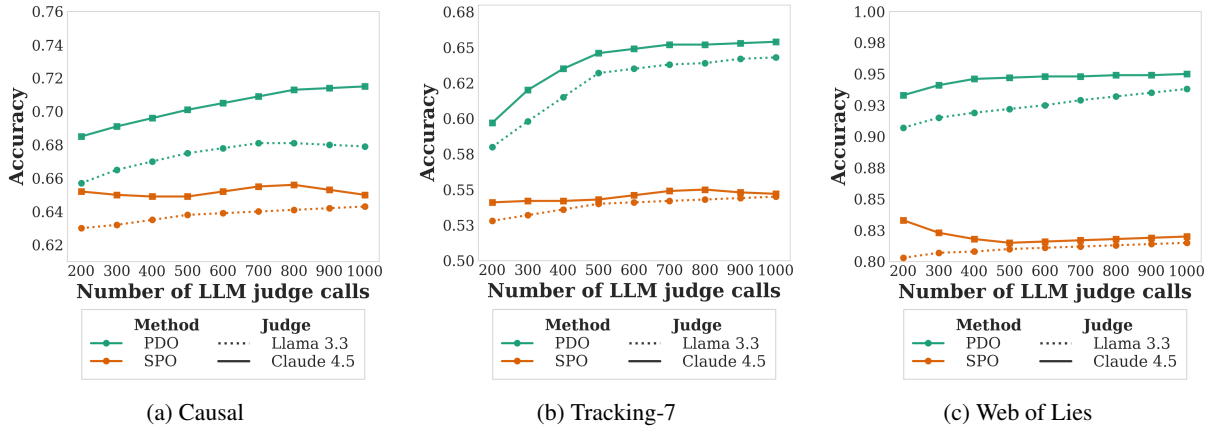


Figure 7: Test performance when using Claude 4.5 vs. Llama 3.3 as the LLM preference judge on BBH tasks under the same judge-call budget. Using Claude 4.5 more reliably identifies higher-performing prompts; the improvement is largest on tasks where the Llama 3.3 judge is noisier (e.g., Causal in Figure 7a and Geometric in Figure 3b).

Dataset	Non-mutate	Mutate	Diff.	<i>p</i>-value
Causal	0.709	0.703	-0.006	> 0.05
Date	0.927	0.938	0.011	0.0291
DisambigQA	0.766	0.760	-0.006	> 0.05
Formal	0.759	0.767	0.008	> 0.05
Geometric	0.611	0.672	0.061	9.7×10^{-5}
Hyperbaton	0.935	0.943	0.008	> 0.05
Logical-5	0.821	0.831	0.010	> 0.05
Logical-7	0.755	0.750	-0.005	> 0.05
Navigate	0.903	0.917	0.014	0.0146
Penguins	0.943	0.955	0.012	0.0236
Salient	0.719	0.711	-0.008	> 0.05
Snarks	0.863	0.876	0.013	> 0.05
Tracking-5	0.793	0.833	0.040	0.0252
Tracking-7	0.585	0.647	0.062	0.0040
Tracking-3	0.949	0.957	0.008	0.0491
Web of Lies	0.919	0.979	0.060	1.5×10^{-6}

Table 7: **Effect of prompt mutation on BBH.** We compare the final Copeland-winner prompt selected by PDO when mutation is disabled (*Non-mutate*) versus when mutation is enabled (*Mutate*). For each BBH task, we report the mean test accuracy over 10 independent runs. We perform a two-sample *t*-test per task to assess whether mutation yields a statistically significant accuracy change. Overall, 8 out of 16 tasks show significant improvement at $p < 0.05$. We visualize representative tasks (Web of Lies and Tracking-7) in Figure 4. Tasks with small negative differences show only minor changes and are likely attributable to run-to-run noise; across all tasks considered, mutation does not materially degrade performance.

D Further Discussion on LLM Judge Design for Closed-Ended Tasks

The design of an LLM preference judge can vary depending on the type of task PDO tackles. In this section, we present the rationale and provide a detailed analysis of our judge design for multiple-choice BBH tasks.

Since outputs in BBH are tied to fixed multiple-choice answers, the judging process naturally splits into two cases. (i) When the two prompts produce different answers, the judge’s role is straightforward: a capable judge should be able to identify the correct answer and select the corresponding prompt. (ii) When the two prompts produce the same answer, the situation is less clear. In this case, the answer alone cannot distinguish between the prompts, so we instead ask the judge to compare their reasoning chains. This design is well suited to BBH tasks, which are explicitly constructed to stress reasoning (Suzgun et al., 2022): the datasets involve multi-step deduction, object tracking, logical rules, or numerical reasoning, where strong performance correlates with the quality of intermediate steps. Our hypothesis is that examining the reasoning trace can reveal whether the model reached the correct answer through a coherent process or by shortcuts or luck.

This formalizes our dual-judge approach, as discussed in Section 5.1, where the judge evaluates either the answer or the reasoning depending on the case. To enable this, we prompt the model to always produce both a reasoning chain and a final answer using JSON-guided decoding.

D.1 Validating the Dual-Judge Approach

To further probe the dual-judge approach, we conduct the following experiments. For each dataset, we select two prompts with different accuracy levels: prompt A as the **higher-accuracy** prompt and prompt B as the **lower-accuracy** prompt. We then test the judge in two scenarios: (i) whether prompt A is preferred when A is correct and B is wrong, and (ii) whether prompt A is preferred when both A and B produce correct answers. The results are reported in Figure 8.

Result Analysis. For most BBH tasks, we find strong evidence that the LLM judge reliably selects the higher-accuracy prompt when A is correct and B is wrong. We also observe moderate evidence that when both prompts produce the same answer,

the judge prefers A based on the quality of its reasoning. Datasets where this preference is consistent under both **Answer** and **Reasoning** conditions (*Hyperbaton*, *Tracking-5*, *Tracking-7*, and *Web of Lies*) show superior performance in Table 1. On the other hand, when the judge is inconsistent in identifying the higher-accuracy prompt A in either condition, performance degrades (*Geometric*, *Salient*).

D.2 Weighted Preference Matrix Update

Motivated by the above analysis in Figure 8, we consider a modification of the PDO algorithm for BBH tasks in which reasoning-based decisions, which are noisier than answer-based ones, are down-weighted so that they contribute less to the preference counts. Let W_{ij} and W_{ji} be the win counts used to form the Beta posteriors $\text{Beta}(W_{ij}+1, W_{ji}+1)$ in D-TS. For a duel decided by source $y \in \{\text{Answer}, \text{Reasoning}\}$, we introduce a margin parameter $\gamma_y \in [0, 0.5]$ and update

$$(W_{ij}, W_{ji}) \leftarrow (W_{ij} + 0.5 + \gamma_y, W_{ji} + 0.5 - \gamma_y)$$

whenever i is preferred, with the symmetric rule when j is preferred. This preserves one effective comparison per duel (sum increment = 1) while shrinking the winner–loser gap to $2\gamma_y$ under noisier judgments. We fix $\gamma_{\text{Answer}} = 0.5$ and ablate $\gamma_{\text{Reasoning}} \in \{0.0, 0.2, 0.5\}$. Figure 9 reports, for each dataset, the ground-truth accuracy *rank* of the Copeland leader over rounds. The results suggest that smaller values of $\gamma_{\text{Reasoning}}$ (e.g., 0.2) often yield more stable progress than the full update $\gamma_{\text{Reasoning}} = 0.5$ on tasks with noisier reasoning judgments, while tasks with consistent reasoning judgments remain robust across settings.

Takeaway. For BBH tasks, we use discounted updates as a simple mechanism to handle differences in judge reliability between answer-based and reasoning-based decisions in preference-based prompt optimization. More generally, for other tasks, LLM judges could be extended with adaptive adjustments that reflect reliability or with additional mechanisms to better capture uncertainty in judgments. Exploring such extensions is an interesting direction for future work.

D.3 Case Study

We conduct a case study of the judge’s decisions in Table 8 on the BBH *Web of Lies* task by comparing

a lower-accuracy Prompt A (0.852) with a higher-accuracy Prompt B (1.000). Two representative scenarios explain why the judge tends to prefer B: (i) when both prompts produce the correct answer, the reasoning-based judge compares the reasoning chains and favors B's more concise and internally consistent rationale; (ii) when the prompts disagree and B's answer is correct, the answer-based judge correctly selects B.




Figure 8: For each BBH dataset, we construct a prompt pair (A, B) , where A is the higher-accuracy prompt and B is the lower-accuracy prompt. We report the proportion of comparisons in which the judge prefers A . **Answer** (green) covers cases where A is correct and B is wrong; the judge evaluates by checking answers. **Reasoning** (blue) covers cases where both prompts produce the correct answer; the judge compares their reasoning chains. *Aggregate* pools all examples across datasets.

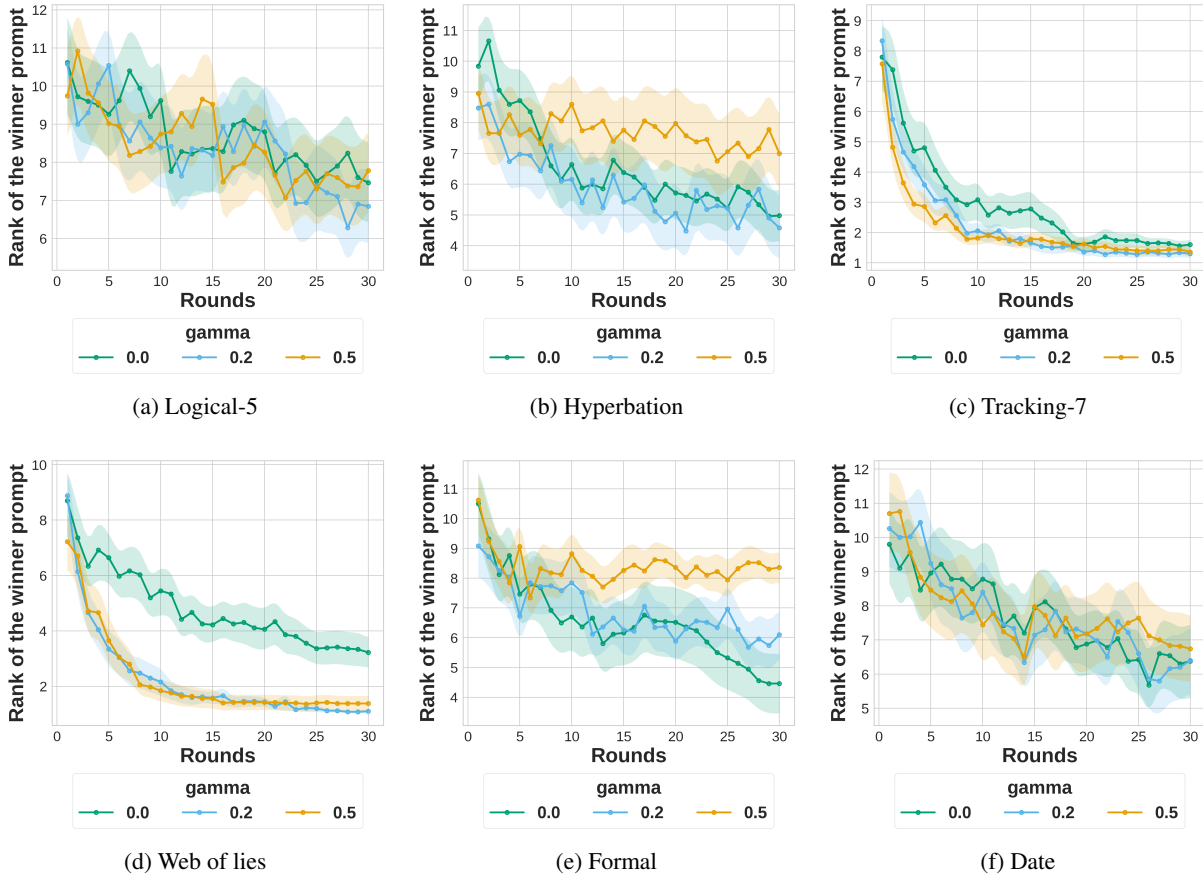


Figure 9: Effect of the reasoning-discount γ in the D-TS update across six BBH datasets. Each plot shows the ground-truth accuracy rank (lower is better) of the current Copeland leader over rounds. We fix $\gamma_{\text{Answer}} = 0.5$ and ablate $\gamma_{\text{Reasoning}} \in \{0.0, 0.2, 0.5\}$. Results indicate that introducing a mild discount at $\gamma_{\text{Reasoning}} = 0.2$ generally accelerates convergence and produces better final ranks overall compared to the undiscounted case $\gamma_{\text{Reasoning}} = 0.5$.

Table 8: Case study of Reasoning-based and Answer-based judging on Prompt A (lower accuracy) and Prompt B (higher accuracy) for selected questions from the *Web of Lies* dataset in BBH.

Item	Content
Prompt definitions	
Prompt A (accuracy = 0.852)	Analyze the given statements about individuals telling the truth or lying, and determine the truthfulness of a specific person by iteratively applying logical deductions, stating assumptions before solving, and considering the chain of statements to arrive at a conclusion.
Prompt B (accuracy = 1.000)	Analyze the given statements about individuals telling the truth or lying, and determine the truthfulness of the person in question by following the chain of statements and applying logical reasoning. For example, if we have the statements: ‘Fidel lies. Fletcher says Fidel tells the truth. Yolanda says Fletcher lies. Raymond says Yolanda lies. Leda says Raymond lies,’ we can conclude that Leda tells the truth by tracing the statements: since Fidel lies, Fletcher’s statement that Fidel tells the truth is false, meaning Yolanda’s statement that Fletcher lies is true, which in turn means Raymond’s statement that Yolanda lies is false, and thus Leda’s statement that Raymond lies is true. Another example is: ‘William tells the truth. Phoebe says William lies. Alejandro says Phoebe lies. Lorine says Alejandro tells the truth. Christie says Lorine tells the truth,’ where we can deduce Christie tells the truth by following the chain of statements in a similar manner. Apply this reasoning process to the given question to determine the truthfulness of the person in question.
Case I: Both prompts output the same correct answer → Reasoning-based judge	
Question	Kristian lies. Sherrie says Kristian lies. Delbert says Sherrie lies. Jerry says Delbert tells the truth. Shalonda says Jerry tells the truth. <i>Does Shalonda tell the truth?</i>
Correct answer	No
Prompt A reasoning	Since Kristian lies and Sherrie says Kristian lies, Sherrie tells the truth. Delbert says Sherrie lies, so Delbert lies, which means Jerry lies when saying Delbert tells the truth, and consequently, Shalonda lies when saying Jerry tells the truth.
Prompt A answer	No
Prompt B reasoning	Since Kristian lies, Sherrie’s statement that Kristian lies is true. This means Delbert’s statement that Sherrie lies is false, so Jerry’s statement that Delbert tells the truth is false, and thus Shalonda’s statement that Jerry tells the truth is false.
Prompt B answer	No
Judge reasoning	Prompt B’s reasoning is more concise and directly addresses the truthfulness of each statement in a straightforward manner. It clearly states the implications of each person’s statement without introducing unnecessary complexity, making it easier to follow and understand. The logical flow from Kristian’s lie to Shalonda’s falsehood is more streamlined in B, enhancing clarity and coherence.
Judge decision	B

Case II: Prompts output *different* answers (B correct, A incorrect)

→ **Answer-based judge**

Question	Vina tells the truth. Helene says Vina lies. Kandi says Helene tells the truth. Jamey says Kandi lies. Ka says Jamey lies. <i>Does Ka tell the truth?</i>
Correct answer	No
Prompt A reasoning	If Ka lies, then Jamey tells the truth, which means Kandi lies, and thus Helene lies, contradicting Vina's truthfulness. Therefore, Ka must tell the truth to maintain the consistency of the statements.
Prompt A answer	Yes
Prompt B reasoning	Since Vina tells the truth, Helene's statement that Vina lies is false. This makes Kandi's statement that Helene tells the truth also false, which in turn makes Jamey's statement that Kandi lies true, and thus Ka's statement that Jamey lies is false, meaning Ka lies.
Prompt B answer	No
Judge reasoning	The correct answer is determined by tracing the statements back to Vina's truthfulness. Response B correctly identifies the chain: Vina true → Helene false → Kandi false → Jamey true → Ka's statement about Jamey false (Ka lies). This matches the task requirement.
Judge decision	B

BBH: Answer-based Preference Judge Template

```
## Role ##
You are a specialized judge focused on evaluating answer correctness when two responses give
different answers.
## Task ##
{question}
## Response from Prompt X ##
**Reasoning:** {reasoning_X}
**Answer:** {answer_X}
## Response from Prompt Y ##
**Reasoning:** {reasoning_Y}
**Answer:** {answer_Y}
## Your Task ##
The responses above give different answers: "{answer_X}" vs "{answer_Y}".
Your job is to determine which answer is more correct for the given task.
## Evaluation Criteria ##
Focus primarily on:
1. **Factual Accuracy** - Which answer better matches reality and task requirements?
2. **Task Alignment** - Which answer better fulfills the specific question asked?
## Output Format ##
{{
  "reasoning": "Your detailed justification explaining why prompt X or Y provided the more
  correct answer (~100 words).",
  "winner": "X or Y"
}}
```

BBH: Reasoning-based Preference Judge Template

```
## Role ##
You are a specialized judge focused on evaluating reasoning quality when two responses give the
same answer.
## Task ##
{question}
## Response from Prompt X ##
**Reasoning:** {reasoning_X}
**Answer:** {answer_X}
## Response from Prompt Y ##
**Reasoning:** {reasoning_Y}
**Answer:** {answer_Y}
## Your Task ##
The responses above give the same answer: "{answer_X}".
Since both arrive at the same conclusion, your job is to determine which reasoning process is
better.
## Evaluation Criteria ##
Focus primarily on:
1. **Logical Coherence** - Is the reasoning chain clear and well-structured?
2. **Completeness** - Does the reasoning address all key aspects of the problem?
3. **Clarity** - Is the reasoning easy to follow and understand?
4. **Accuracy** - Are the intermediate steps and assumptions correct?
## Output Format ##
{{
  "reasoning": "Your detailed justification explaining why prompt X or Y provided better
  reasoning (~100 words).",
  "winner": "X or Y"
}}
```

MS-MARCO: Preference Judge Template

```
## Role ##
You are a meticulous, impartial referee evaluating two competing answers to determine which
better answers the given question based on the provided context.
## Query ##
{query}
## Context ##
{context}
## Answer from Prompt X ##
{answer_X}
## Answer from Prompt Y ##
{answer_Y}
## Evaluation Criteria ##
Compare both answers based on:
1. Accuracy - How factually correct is each answer based on the context?
2. Completeness - Does the answer address all aspects of the question?
3. Relevance - How well does the answer stay focused on answering the question?
4. Clarity - How clear and well-articulated is the answer?
## Output Format ##
{{
  "reasoning": "Your detailed justification explaining why answer X or Y is better (~100 words)
  .",
  "winner": "X or Y"
}}
```

MS-MARCO: final evaluations with ground-truth references

```
"""
Begin your evaluation by carefully comparing the AI-generated answer with the reference solution.
Identify any outputs that are not true, as well as omissions or deviations, and simulate
human judgments in explaining how these impact the overall quality of the response. Ensure
that your assessment is objective and consistent.
At the end of your evaluation, assign a score from 1 to 5 based on the following scale:

- 1: Very poor - does not meet the requirements or is significantly incorrect.
- 2: Poor - contains major errors or omissions.
- 3: Fair - adequate but with notable flaws.
- 4: Good - meets the requirements with minor errors.
- 5: Excellent - fully accurate and well-articulated.

[User Question]
{question}

[Reference Solution]
{ground_truth}

[AI-Generated Answer]
{prediction}

Your response should be a valid JSON string (no backticks) following this schema:
{{ "explanation": "{{Detailed reasoning based on the comparison}}"
  "score": {{1-5}}
}}
"""
```

Template for generating initial prompts

```
"""
You are an expert prompt-engineer. Your task is to generate exactly 1 high-quality system-level instruction for the target reasoning task.

# Dataset Snapshot
Below is an LM-written summary of the unlabeled question pool:
{dataset_summary}

# Sample Inputs (do NOT answer them)
{questions}

# Prompt-Engineering Tip
{tip}

# Output Format (STRICT)
Return exactly 1 item in a JSON array, and nothing else:
[ "Your single high-quality instruction here" ]
"""
```

Template for mutating on top-performing prompts

```
"""
You are an expert prompt-engineer specializing in prompt optimization.
Your task is to generate 1 diverse, high-quality mutation of the currently BEST PERFORMING instruction for the target reasoning task.

# BEST PERFORMING Instruction (Current Champion)
{instructions}

# CRITICAL: Follow This Prompt-Engineering Tip
{tip}

# Output Format
Return exactly 1 mutated instruction in a JSON object, and nothing else:
{ "mutated_prompt": "Your mutated instruction here, following the tip." }
"""
```

Prompt engineering tips for prompt generations

```
# Mutation tips for top-performer-guided prompt mutations
MUTATION_TIPS = {
    "expansion": "Keep the current champion instruction exactly as is, but expand on it by adding
        additional helpful guidance or clarifications. The result should be the original
        instruction plus new supplementary content.",
    "minimal": "Make very minimal changes to the current champion instruction. Keep it around the
        same length and modify only a few words through paraphrasing while preserving the core
        meaning.",
    "few_shot": "Add a few concrete examples to the current champion instruction to demonstrate
        the expected reasoning process or output format. Include 1-3 brief example cases that show
        how to apply the instruction.",
    "emphasis": "Adjust the tone, emphasis, or directional focus of the current champion
        instruction to create different reasoning patterns.",
}

# Initial instruction generation tips
INITIAL_INSTRUCTION_TIPS = {
    "framing": "Set the context for the task by framing it as a concrete creative scenario.",
    "simple": "Keep the instruction clear and concise.",
    "description": "Make sure your instruction is very informative and descriptive.",
    "persona": "Provide the LM with a creative persona that is relevant to the task.",
    "edge_cases": "List tricky cases the instruction must handle.",
    "assumptions": "Have the model state assumptions before solving.",
}
```