

CE-RM: A Pointwise Generative Reward Model Optimized via Two-Stage Rollout and Unified Criteria

Xinyu Hu^{1,3}, Yancheng He³, Weixun Wang³, Tao Feng³, Li Lin¹,
Jiashun Liu^{2,3}, Wenbo Su³, Bo Zheng³, Xiaojun Wan¹

¹Wangxuan Institute of Computer Technology, Peking University

²Hong Kong University of Science and Technology ³Alibaba Group

{huxinyu,wanxiaojun}@pku.edu.cn, {heyancheng.hyc,weixun.www}@taobao.com

Abstract

Automatic evaluation is crucial yet challenging for open-ended natural language generation, especially when rule-based metrics are infeasible. Compared with traditional methods, the recent LLM-as-a-Judge paradigms enable better and more flexible evaluation, and show promise as generative reward models for reinforcement learning. However, prior work has revealed a notable gap between their seemingly impressive benchmark performance and actual effectiveness in RL practice. We attribute this issue to some limitations in existing studies, including the dominance of pairwise evaluation and inadequate optimization of evaluation criteria. Therefore, we propose **CE-RM-4B**, a pointwise generative reward model trained with a dedicated two-stage rollout method, and adopting unified query-based criteria. Using only about 5.7K high-quality data curated from the open-source preference dataset, our CE-RM-4B achieves superior performance on diverse reward model benchmarks, especially in Best-of-N scenarios, and delivers more effective improvements in downstream RL practice.

1 Introduction

Automatic evaluation has long been a widely studied yet challenging area in natural language processing (Papineni et al., 2002; Yuan et al., 2021; Gao et al., 2025). Although verifiable tasks such as question answering and mathematics can be evaluated using rule-based metrics, they still face practical difficulties, including diverse answer formats and reasoning process validation (Lightman et al., 2023; Shao et al., 2024a). Moreover, there exists a large number of open-ended generation tasks with no well-defined correct answers, such as creative writing and general instruction following (Lou et al., 2024). Therefore, it is crucial to develop flexible and efficient automatic evaluation methods.

Traditional research on automatic evaluation has generally evolved through two stages: string-matching metrics such as BLEU (Papineni et al., 2002) and ROUGE (Lin, 2004), and language-model-based metrics such as BARTScore (Yuan et al., 2021) and COMET (Rei et al., 2020). Despite the substantial improvements in evaluation performance, these metrics largely rely on high-quality references and exhibit poor robustness, which can be easily affected by semantic perturbations (He et al., 2023). With the recent rapid development of large language models (LLMs), their improving instruction-following capability has driven a new LLM-based automatic evaluation paradigm (LLM-as-a-Judge) (Wang et al., 2023; Gao et al., 2025), which formulates evaluation as an end-to-end generation task to imitate human judgment and supports diverse evaluation functionalities.

Compared with traditional evaluation metrics, this new paradigm achieves performance closer to human-level evaluation. More importantly, it can serve as a generative reward model (GRM) (Mahan et al., 2024) to provide automatic reward signals for reinforcement learning (RL), especially in scenarios that are difficult to verify with rules. Existing studies (Wang et al., 2024; Liu et al., 2025b; Chen et al., 2025b,a; Guo et al., 2025; Whitehouse et al., 2025) have made substantial efforts to enhance the performance of GRMs, including distilling data from state-of-the-art LLMs, strengthening the reasoning process during evaluation, and employing rejection sampling for self-bootstrapping. In addition, some work (Liu et al., 2024b; Xu et al., 2024; Liu et al., 2025b) has taken notice of a crucial component of evaluation—criteria—and demonstrated that basing evaluation on high-quality criteria can significantly enhance performance.

Despite the increasing performance of these studies on reward model benchmarks, Wen et al. (2025) have shown that there remains a notable gap between such benchmark performance and actual ef-

fectiveness when deployed in RL. We argue that one primary reason may be that most GRMs adopt a **pairwise** evaluation protocol. In fact, existing public preference training data are almost pairwise: they are relatively easy to construct and can be generated automatically at scale. In contrast, **pointwise** data require high-quality human annotations and are thus more expensive. Furthermore, many RM benchmarks are also designed as pairwise comparisons. However, this pairwise protocol is difficult to leverage in RL practice. For example, to apply such RMs in RL with GRPO (Shao et al., 2024b), one often needs to compare each pair of trajectories within each rollout group and then use algorithms like Elo (Elo and Sloan, 1978) to compute a pointwise score for each trajectory, leading to quadratic computational cost and low efficiency.

In addition, although some studies have highlighted the importance of evaluation criteria, they typically treat criteria as part of the evaluation process and optimize the entire evaluation holistically, rather than explicitly strengthening the criteria generation. More notably, their criteria are often generated jointly conditioned on both the query and response. While this approach may yield more comprehensive criteria by incorporating the query requirements and the specific response content, it can lead to inconsistent criteria when multiple responses to the same query are evaluated and compared, thereby introducing biases. This setting is, however, closer to real-world scenarios, such as computing rewards within a rollout group. In Section 2, we conduct experiments demonstrating that using unified criteria solely based on the query is increasingly beneficial as the number of responses per query grows.

Motivated by these considerations, we propose a reinforcement learning method with a two-stage rollout—unified Criteria then Evaluation—to train our pointwise GRM, **CE-RM-4B**. Specifically, our model conducts evaluation by first generating evaluation criteria solely based on the query, and then, conditioned on the query and the criteria, producing a detailed evaluation analysis for each response. To train the model, we select about 5.7K data from the public pairwise preference dataset Skywork-Reward-Preference-80K-v0.2 (Liu et al., 2024a) through a series of dedicated quality-filtering procedures. During reinforcement learning, we further design specific algorithms to leverage the original pairwise preference labels to estimate fine-grained reward signals for criteria and evaluation trajectory

ries, respectively, without requiring any additional pointwise score annotations.

We conduct experiments on multiple commonly used RM benchmarks, covering diverse and challenging tasks as well as more realistic Best-of-N scenarios. Experimental results show that, compared with existing work, our CE-RM-4B achieves superior performance with only 4B parameters and less training data. Moreover, on benchmarks where each query is associated with more responses, our model indeed exhibits stronger evaluation capability. We also perform extensive ablation analyses to verify the importance of explicitly optimizing criteria generation and the effectiveness of our proposed methods. Furthermore, we apply our CE-RM-4B in practical RL with GRPO, and observe consistently better improvement. In summary, our contributions are as follows:

1. We reveal several limitations of current generative reward models, including the dominance of pairwise evaluation and the lack of dedicated optimization for evaluation criteria.
2. We propose a two-stage rollout reinforcement learning method that unifies evaluation criteria conditioned on the query and effectively utilizes pairwise preference labels to train our pointwise GRM, CE-RM-4B.
3. Our model achieves superior performance on common reward model benchmarks, especially in more realistic Best-of-N scenarios, and yields more effective improvements in RL practice. Our model is available at huggingface.co/PKU-ONELab/CE-RM-4B.

2 Preliminary Study

To investigate the role of evaluation criteria and how their different usages affect evaluation performance, we conducted preliminary experiments under three pointwise evaluation settings. For each query x and its candidate responses $\{y_i\}$, the first setting follows the most common practice: the model performs the analysis and then outputs a score (*Direct Evaluation*). It can be formalized as follows, where r_θ denotes the GRM and e_i denotes the evaluation output, f_{score} is a rule-based function that extracts the score s_i from the evaluation.

$$e_i \sim r_\theta(x, y_i), \quad s_i = f_{\text{score}}(e_i) \quad (1)$$

In this setting, we do not explicitly require the model to provide evaluation criteria. Nevertheless,

Evaluation Setting	RWBench	RWBench2	RM-Bench
Direct Evaluation	90.0	76.7	79.8
Explicit Criteria	90.3	78.8	81.5
Unified Criteria	90.6	79.8	83.0

Table 1: The results of Qwen3-Max with three different evaluation settings on three reward model benchmarks: RewardBench (RWBench), RewardBench2 (RWBench2) and RM-Bench.

the model is allowed to benefit from its reasoning process, and in many cases, it implicitly constructs certain criteria.

The second setting requires the model to first generate explicit evaluation criteria and then perform corresponding analysis and scoring, conditioned on both the query and the response (*Explicit Criteria*), where c_i denotes the generated criteria.

$$c_i, e_i \sim r_\theta(x, y_i), s_i = f_{\text{score}}(e_i) \quad (2)$$

The third setting asks the model to first generate the evaluation criteria solely based on the query, and then evaluate each response conditioned on the query and this **unified criteria** c (*Unified Criteria*), as formalized below.

$$c \sim r_\theta(x), e_i \sim r_\theta(x, y_i, c), s_i = f_{\text{score}}(e_i) \quad (3)$$

We evaluate them on three widely used RM benchmarks: RewardBench (Lambert et al., 2025b), RewardBench2 (Malik et al., 2025), and RM-Bench (Liu et al., 2025a), where each query is associated with 2, 4, and 6 responses, respectively. More descriptions are provided in Section 4.2. We use Qwen3-Max (Team, 2025b) as the GRM, and the prompts for the three settings are presented in Figures 4–6 in Appendix A. As shown in Table 1, explicitly generating criteria consistently improves evaluation performance. Furthermore, as the number of responses per query increases, the advantage of unified criteria becomes more pronounced.

3 Methodology

Motivated by the above exploration, we aim to train a pointwise GRM with unified criteria, whose evaluation process follows Eq. (3), while explicitly strengthening the criteria generation capability. To this end, we propose a specific reinforcement learning method based on a two-stage rollout. In addition, we design algorithms to convert pairwise preference labels into fine-grained rewards for criteria and evaluation trajectories, respectively.

3.1 Data Curation

To obtain training data with high quality and instance efficiency, we start from a widely used preference dataset, Skywork-Reward-Preference-80K-v0.2, and perform data filtering. Each instance contains a query and a pair of chosen and rejected responses (x, y^c, y^r) , corresponding to a pairwise preference label. We first use Qwen3-4B-Instruct-2507 (Team, 2025c) to perform multiple pointwise evaluations for each instance, and compute accuracy as the fraction of trials in which the chosen response y^c is assigned a higher score than the rejected response y^r . We then retain about 29K instances whose accuracy does not exceed 0.6, representing those the model is uncertain about or judges incorrectly. Furthermore, to ensure the diversity of queries, we employ Qwen3-Max to identify the task type of each query and cluster them using Qwen3-4B-embedding (Zhang et al., 2025c). Finally, we perform stratified random sampling to obtain approximately 12K instances whose task-type distribution is as balanced as possible.

3.2 Cold Start

We distill a small amount of supervised fine-tuning (SFT) data from Qwen3-Max for cold-starting our reward model under the two-stage evaluation protocol. For each query x , following the prompt in Figure 6 and Eq. (3), we require Qwen3-Max to first generate three sets of criteria $\{c_i\}_{i=1}^3$ solely based on the query. Then, conditioned on x and each c_i , the model generates three evaluations for the chosen response y^c and the rejected response y^r , denoted as $\{e_{ij}^c\}_{j=1}^3$ and $\{e_{ij}^r\}_{j=1}^3$, respectively.

In particular, the evaluation is conducted based on each criterion item in the generated criteria set with the corresponding sub-score, and then the overall analysis and final score are derived based on these contents to ensure reliability. Moreover, to prevent the criteria set from failing to cover specific aspects of a response, we allow the model to introduce an additional adjustment beyond the criteria when necessary, as shown in stage 2 of Figure 6. The score range is set to 0–10 with a half-point increment to mitigate the limited granularity of pointwise scoring. In preliminary tests, this setting achieves a good balance between granularity and reliability, while finer granularity, such as 0.1, often causes the scores to deviate from well-grounded evidence.

Since each instance only contains the preference

label and no human evaluation as a gold standard, we design empirical measures to assess the quality of the instance, criteria, and evaluation, making the cold-start data as reliable as possible.

Instance We select the instance where, based on all three criteria sets, the corresponding evaluations consistently yield the higher scores for the chosen response than the rejected response, ensuring the correctness of the preference.

$$s_{ij}^c = f_{\text{score}}(e_{ij}^c), \quad s_{ik}^r = f_{\text{score}}(e_{ik}^r) \quad (4)$$

$$\prod_{1 \leq i \leq 3} \prod_{1 \leq j, k \leq 3} \mathbb{I}(s_{ij}^c > s_{ik}^r) = 1 \quad (5)$$

Criteria For each retained instance above, among its three criteria sets, we select e_{i^*} where the variance of its induced evaluation scores is the smallest. If this variance exceeds a threshold, the instance is discarded. This aims to ensure that the criteria are unambiguous, enabling stable and correct scoring during evaluation.

$$i^* = \arg \min_{1 \leq i \leq 3} \left(\text{Var}(\{s_{ij}^c\}_{j=1}^3) + \text{Var}(\{s_{ij}^r\}_{j=1}^3) \right) \quad (6)$$

Evaluation For each selected instance and its criteria set, we choose the response evaluation whose score is the median among the three. Furthermore, to reuse these data in subsequent reinforcement learning while preventing the model from directly observing the preference label, we keep only one response (either chosen or rejected) along with its corresponding evaluation. And we design a specific algorithm for the aforementioned response selection to make the resulting score distribution as uniform as possible.

Finally, we obtain about 2.2K instances, denoted as $\mathcal{D}_{\text{SFT}} = \{(x_n, y_n, c_n, e_n)\}_{n=1}^N$, where each contains a query, a response, a criteria set, and an evaluation. The cold-start training objective is defined as follows. For each instance, we simultaneously optimize the generation of criteria and evaluation, but under different conditioning contexts.

$$\mathcal{L}_{\text{SFT}} = -\frac{1}{N} \sum_{n=1}^N \left(\sum_t \log r_\theta(c_{n,t} | x_n, c_{n,<t}) + \sum_t \log r_\theta(e_{n,t} | x_n, y_n, c_n, e_{n,<t}) \right) \quad (7)$$

3.3 RL with Two-Stage Rollout

Unlike the cold start, reinforcement learning requires data in the original triplet form (x, y^c, y^r) . Since criteria and evaluations are not needed, we relax the filtering condition used in cold start, and keep instances where at least one set of criteria leads to entirely correct evaluations as follows, resulting in about 5.7K instances, denoted as \mathcal{D}_{RL} .

$$\sum_{1 \leq i \leq 3} \prod_{1 \leq j, k \leq 3} \mathbb{I}(s_{ij}^c > s_{ik}^r) \geq 1 \quad (8)$$

The evaluation process conducted by our reward model consists of two stages: criteria generation and evaluation generation, both of which need to be optimized. Therefore, building upon the reinforcement learning with GRPO (Shao et al., 2024b), we adopt a two-stage rollout scheme. Similar to the procedure described in Section 3.2, for each query, the model first generates n_c criteria trajectories in the first stage; then, conditioned on the query and each generated criteria trajectory, it produces n_e evaluation trajectories for the chosen and rejected responses, respectively, in the second stage.

Given that only pairwise preference labels are available, we design the following methods to estimate fine-grained reward signals for criteria and evaluation trajectories, improving both reliability and discriminability. The reward for each set of criteria c_i is defined as the win rate of pairwise comparison between the chosen and rejected responses in terms of their induced scores from the subsequent corresponding evaluations, estimated in an aggregated manner to reduce bias.

$$\mathcal{R}_{\text{criteria}} = \frac{1}{n_e^2} \sum_{1 \leq j, k \leq n_e} \mathbb{I}(s_{ij}^c > s_{ik}^r) \quad (9)$$

The reward for each evaluation e_{ij}^c of the chosen response y^c is defined as its win rate in pairwise score comparisons against all evaluations of the rejected response, multiplied by a binary format reward (0 for incorrect, 1 otherwise); and the reward for each evaluation e_{ij}^r of y^r is defined analogously.

$$\mathcal{R}_{\text{eval}}^c = \frac{1}{n_e} \sum_{1 \leq k \leq n_e} \mathbb{I}(s_{ij}^c > s_{ik}^r) \cdot \mathcal{R}_{\text{format}} \quad (10)$$

Notably, in this two-stage rollout, the criteria and the evaluations of chosen and rejected responses are all treated as optimization objectives for the current training instance, but they correspond to three different conditional generations. Therefore,

Reward Model	Data	RWBench	RWBench2	RM-Bench	PPE Corr	JudgeBench	Overall
<i>Pairwise RM</i>							
GPT-4o (Hurst et al., 2024)	-	86.7	-	72.5	57.6	56.6	-
Llama-3.3-70B-Instruct	-	85.4	-	69.5	65.7	48.6	-
JudgeLRM-7B (Chen et al., 2025a)	100K	75.2	-	58.7*	42.6*	54.6*	-
RM-R1-Qwen-7B (Chen et al., 2025b)	73K	85.2	-	70.2	57.9*	59.1*	-
RM-R1-Qwen-32B (Chen et al., 2025b)	73K	91.4	-	79.1	59.3*	68.6*	-
RRM-7B (Guo et al., 2025)	420K	82.2	-	70.8*	60.1*	60.9*	-
RRM-32B (Guo et al., 2025)	420K	91.2	-	<u>82.8*</u>	67.9*	<u>75.1*</u>	-
RISE-Judge-7B (Yu et al., 2025)	73K	88.2	-	68.0*	60.4*	59.7*	-
RISE-Judge-32B (Yu et al., 2025)	40K	92.7	-	72.3*	55.0*	66.3*	-
Think-J-32B (Huang et al., 2025)	9.8K	90.5	-	79.8	-	-	-
CompassJuder2-7B (Zhang et al., 2025b)	?	91.0	-	71.5*	60.2*	63.1	-
CompassJuder2-32B (Zhang et al., 2025b)	?	92.6	-	73.2*	56.6*	65.5	-
J1-Llama-8B (Whitehouse et al., 2025)	22K	85.7	-	73.4	59.2	42.0	-
J1-Llama-70B (Whitehouse et al., 2025)	22K	<u>93.3</u>	-	82.7	<u>72.9</u>	60.0	-
<i>Pointwise RM</i>							
Gemini-2.5-Flash (Team, 2025a)	-	80.7	<u>75.9</u>	70.8	61.9	66.9	71.2
CompassJuder1-32B (Cao et al., 2024)	900K	81.2*	56.5*	54.4*	48.0*	42.3*	56.5*
DeepSeek-GRM-27B (Liu et al., 2025b)	237K	86.0	-	-	59.8	-	-
CLoud-Gemma-2-27B (Liu et al., 2025b)	237K	82.0	-	-	62.4	-	-
J1-Llama-8B (Whitehouse et al., 2025)	22K	-	-	-	53.8	57.7	-
J1-Llama-70B (Whitehouse et al., 2025)	22K	-	-	-	65.0	60.0	-
TIR-Judge-Distill-4B (Xu et al., 2025)	26K	76.6	67.3	71.9	65.9	66.7	69.7
TIR-Judge-Zero-4B (Xu et al., 2025)	26K	77.3	68.3	72.8	69.8	70.4	71.7
TIR-Judge-Distill-8B (Xu et al., 2025)	26K	81.0	71.6	76.7	71.0	68.2	73.7
TIR-Judge-Zero-8B (Xu et al., 2025)	26K	81.4	73.4	76.3	70.3	67.5	73.8
CE-RM-4B (Ours)	5.7K	89.0	74.6	79.8	69.7	73.7	<u>77.4</u>
CE-RM-4B (Scaling@2)	5.7K	89.4	75.9	81.9	73.1	75.7	79.2
CE-RM-4B (Scaling@4)	5.7K	90.0	76.3	83.2	75.0	76.3	80.2

Table 2: The results of our CE-RM-4B compared with other baselines on five common reward model benchmarks, as well as their overall performance. RWBench, RWBench2, and PPE Corr represent the RewardBench, RewardBench2, and PPE Correctness, respectively. And scaling@k denotes the performance obtained by applying test-time scaling for k times and using the averaged evaluation scores. The results marked with * are those not reported in the original paper but tested and obtained by ourselves; The results highlighted in **bold** and underlining indicate the best one in each column, and the latter refers specifically to comparisons conducted without test-time scaling.

unlike typical GRPO, which groups rollouts by instance, we partition trajectories into three sub-groups (criteria sets, chosen-response evaluations, rejected-response evaluations) for subsequent relative advantage estimation and optimization.

4 Experiments

4.1 Implementation Details

We use Qwen3-4B-Instruct-2507 (Team, 2025c) as the foundation model for training, to balance the model capability and experimental cost. For the cold start, we set the batch size to 64, the learning rate to 1e-5, and train for 3 epochs. For reinforcement learning, the mini-batch size is 64, with 4 mini-batches per step; the learning rate is 2e-6, and the KL coefficient is 1e-3. More concretely, for

each query, we generate four sets of criteria in the first-stage rollout ($n_c = 4$); in the second stage, we generate two evaluations for both the chosen and rejected responses ($n_r = 2$). Consequently, each instance involves 4 criteria trajectories and 16 evaluation trajectories, i.e., 20 trajectories in total. Moreover, our reinforcement learning is implemented based on the open-source ROLL framework (Wang et al., 2025b). During inference, we set the temperature to 0 to improve determinism.

4.2 Benchmarks

To comprehensively evaluate the performance of our model, we adopt a collection of widely used reward model benchmarks, covering diverse task types, challenging settings, and Best-of-N scenar-

Trajectory	Criteria (n_c)	Evaluation (n_e)	Overall
9	1	4	72.7
10	2	2	74.1
12	4	1	73.6
18	2	4	74.4
21	3	3	75.3
20	4	2	77.4

Table 3: The overall performance of the reward model trained with different combinations of two-stage rollout numbers on five benchmarks. The columns of trajectory, criteria, and evaluation represent the number of total trajectories, generated criteria trajectories and evaluation trajectories for each instance during RL, respectively.

ios. Specifically, we select:

RewardBench (Lambert et al., 2025b) One of the earliest and most popular benchmarks for reward models, spanning multiple categories such as chat, safety, and reasoning. However, its difficulty may be relatively low for current models.

RewardBench2 (Malik et al., 2025) A substantial upgrade in both difficulty and diversity compared with RewardBench. It features more realistic human queries and a global Best-of-N evaluation protocol, and has been shown to correlate more strongly with downstream RLHF performance.

RM-Bench (Liu et al., 2025a) Focuses on robustness under subtle content differences and stylistic biases, testing whether a reward model truly judges based on content rather than style.

PPE Correctness (Frick et al., 2025) Built upon verifiable and challenging benchmarks (MATH, MBPP, MMLU-Pro, GPQA, IFEval) to test alignment in Best-of-N tasks. It has likewise been validated to reflect RLHF performance.

JudgeBench (Tan et al., 2025) Derived from a set of real-world and difficult evaluation datasets (e.g., MMLU-Pro, LiveBench). It assesses the capability to discriminate answer correctness through pairwise comparison.

4.3 Main Results

In Table 2, we compare our CE-RM-4B with a great range of latest generative reward models and several general-purpose LLMs, which are grouped into two categories by evaluation protocol: pairwise and pointwise. Since many prior works do not report results on all benchmarks, we test those




Figure 1: The variance of the evaluation scores corresponding to the same criteria trajectory (in-criteria), and to all criteria trajectories (all-criteria) for each instance, as well as the criteria rewards during RL.

open-source reward models ourselves and supplement the results, which are marked with * in the table. Moreover, as RewardBench2 is not compatible with pairwise reward models by design, corresponding results are unavailable.

The experimental results show that our CE-RM-4B achieves superior overall evaluation performance, despite using much less training data and a smaller model size than other baselines. Its advantage is particularly pronounced on RewardBench2 and PPE Correctness, which are aligned with Best-of-N scenarios, as well as on RM-Bench, where each query is associated with a larger number of responses. It highlights the benefit of our enforcing unified query-based evaluation criteria.

4.3.1 Test-Time Scaling

Similar to work such as Liu et al. (2025b), our model can be improved via test-time scaling using averaged scores from multiple inferences. Since pointwise scores from GRMs often have limited granularity, ties occur commonly and can degrade evaluation performance. And averaging multiple sampled scores effectively mitigates this issue. Table 2 also reports results of scaling at 2 and 4 times. The gains are most evident on RM-Bench, PPE Correctness, and JudgeBench, where ties are not explicitly handled. Moreover, owing to its relatively fewer parameters, our model incurs low computational overhead for test-time scaling, and it even outperforms other substantially larger reward models that employ pairwise evaluation protocols.

4.3.2 Two-Stage Rollout Combination

We further explore combinations of trajectory numbers in our two-stage rollout method. As shown




Figure 2: The results of the ablation experiments that compare our CE-RM-4B with other reward models trained without different key components of our methods.

in Table 3, increasing total trajectories for each instance generally leads to better performance. Under comparable trajectory budgets, adopting 4 criteria trajectories and 2 evaluation trajectories yields the best result, suggesting that while maintaining a certain number of evaluations to stabilize the estimation of reward scores, optimizing the criteria is more critical, which confirms our motivation.

4.3.3 Evaluation Robustness

In Figure 1, we show the variance of scores from the evaluations based on the same criteria trajectory (in-criteria), and on all criteria trajectories (all-criteria) for each instance during RL. Benefiting from the effective optimization of both the criteria and evaluation, our model produces increasingly stable evaluation scores as training proceeds. Similar results are also observed on the test benchmarks, as shown by the comparisons in Table 4, which improves the robustness of CE-RM-4B in practical deployments, such as the low-cost or low-latency settings where only a single inference is allowed.

5 Ablation Analyses

We further conduct ablation studies on key components of our method to validate their importance, whose training settings are similar to Section 4.1:

Unified Criteria This is the most critical component and one of our starting points. For comparison, we train another reward model using the traditional scheme that generates criteria and evaluation jointly conditioned on the query and response (i.e., Eq. (2)), with a similar data construction and total trajectory number during the RL procedure.

Model	Intra	Inter	Acc
Qwen3-4B-Instruct-2507	0.621	1.423	72.8
+ Cold Start	0.482	0.893	76.5
CE-RM-4B w/o Optimized Criteria	0.227	0.515	80.0
CE-RM-4B (Ours)	0.181	0.441	81.1

Table 4: The score variance and accuracy of different models averaged across RewardBench, RewardBench2, and RM-Bench. Intra and Inter denote the variance of the evaluation scores across four samplings, where the former is computed from four evaluations generated based on one sampled criteria set, whereas the latter is calculated by sampling four different sets of criteria and the corresponding evaluations respectively. Specifically, CE-RM-4B w/o Optimized Criteria refers to a compositional setting where the criteria are generated by the RM trained only with the cold start, and the evaluations are then produced based on them by our CE-RM-4B.

Criteria Optimization In the two-stage rollout during reinforcement learning, after computing the reward signals, we discard criteria trajectories in the optimization objective to verify the importance of explicitly enhancing criteria generation.

Evaluation Optimization Similarly, we discard evaluation trajectories in the optimization objective, to examine the influence of optimizing only the criteria.

Sub-Grouping We replace our sub-grouping strategy with the conventional grouping over all trajectories within the rollout of each instance.

Data Quality Filtering We significantly simplify the quality filtering pipelines used in both cold start and RL, retaining only basic correctness checks on

overall judgments, to demonstrate the necessity of our carefully designed quality estimation.

Reinforcement Learning We also evaluate the model trained only with the cold start to quantify the improvements brought by RL.

Figure 2 presents the results of the ablation experiments above, highlighting the effectiveness of the unified query-based criteria and our proposed training methodology.

6 RL Practice

Since most generative reward models adopt the pairwise evaluation, their downstream RLHF performance is typically validated by preference label annotation and DPO-style training (Rafailov et al., 2023). To better approximate real-world scenarios, we instead adopt on-policy reinforcement learning, using the evaluation scores from RMs as the reward signal to optimize the policy model with GRPO. In particular, we randomly sample 50K queries from Tulu 3 dataset (Lambert et al., 2025a) as training data. And we evaluate the effectiveness of policy optimization using Arena Hard (Li et al., 2024a,b), a popular automatic evaluation arena that contains challenging real-world user queries and employs powerful LLMs as the judge, showing strong correlation with human judgments.

Given that most pointwise GRMs are either not open-sourced or perform poorly, we take as baselines the larger-parameter CompassJuderger1-32B and the reward model trained without unified criteria as described in Section 5. They are compared against our CE-RM-4B, as well as that using test-time scaling at four times. We choose Qwen3-8B as the policy model to be optimized, and all experiments apply the similar training settings as those in Section 4.1 with a total of 250 training steps. We also examine the impact of different group sizes in GRPO on optimization performance. As shown in Table 5 and Figure 3, our model yields stable improvements during RL and significantly outperforms the baseline RMs. Furthermore, both increasing the number of rollout trajectories and applying test-time scaling bring substantial additional gains, enabling the optimized policy model to surpass Qwen3-14B.

7 Related Work

The reward model is an important research direction in automatic evaluation, as it provides critical




Figure 3: During training, the evolution of performance of the policy model across three evaluation scenarios, along with the averaged evaluation scores that the reward model assigns to the policy rollouts. In the RL practice presented here, the policy model is Qwen3-8B and the reward model is our CE-RM-4B (without test-time scaling), with GRPO and a group size of 8.

feedback for model alignment. Traditional discriminative reward models are typically trained based on the Bradley–Terry model (Bradley and Terry, 1952), and have been shown to suffer from weak calibration and cross-prompt generalization (Sun et al., 2025; Zhang et al., 2025a). With the rapid development of LLMs, LLM-based generative reward models (GRMs) have emerged, which replace a single scalar score with natural language evaluations. This enables more flexible and informative rewards with better interpretability. In addition, they can leverage test-time scaling to further improve evaluation robustness (Liu et al., 2025b).

A major challenge in developing GRMs lies in obtaining high-quality supervised evaluation data. Building on existing preference datasets, a common approach is to distill chains of thought from frontier LLMs to enrich the original preference labels (Chen et al., 2025b; Huang et al., 2025; Guo et al., 2025). Another direction allows the target model itself to generate candidate evaluations, which are then checked with preference labels to perform rejection finetuning in a self-bootstrapping manner (Wang et al., 2024; Whitehouse et al., 2025). Other studies improve the generalization ability of GRMs by mixing diverse evaluation tasks, including even general instruction-following data, into training (Cao et al., 2024; Wang et al., 2025a).

In addition, since most preference datasets and generative reward models adopt a pairwise comparison protocol, the correctness of the model’s evaluation judgment can be easily verified. Early work

Model	Arena-Hard v0.1		v2 Hard Prompt		v2 Creative Writing	
	Score	CI	Score	CI	Score	CI
Qwen3-8B	66.5	(-2.7 / +2.1)	9.8	(-1.0 / +0.8)	14.4	(-1.4 / +1.8)
Qwen3-14B	77.4	(-2.6 / +1.4)	17.1	(-1.8 / +1.4)	24.7	(-2.1 / +2.0)
<i>GRPO with group size = 4</i>						
Qwen3-8B + CompassJuder1-32B	75.0	(-2.2 / +2.0)	13.4	(-0.8 / +1.1)	27.1	(-2.6 / +2.3)
Qwen3-8B + RM w/o unified criteria	71.0	(-2.5 / +2.0)	12.9	(-1.2 / +1.0)	37.6	(-2.8 / +2.5)
Qwen3-8B + CE-RM-4B (Ours)	75.7	(-1.5 / +2.0)	16.3	(-1.4 / +1.3)	42.9	(-2.2 / +2.3)
Qwen3-8B + CE-RM-4B (Scaling@4)	78.3	(-1.6 / +1.4)	17.6	(-1.2 / +1.4)	48.0	(-2.2 / +2.1)
<i>GRPO with group size = 8</i>						
Qwen3-8B + CompassJuder1-32B	74.7	(-1.8 / +1.8)	13.6	(-1.0 / +1.1)	39.4	(-2.1 / +2.4)
Qwen3-8B + RM w/o unified criteria	72.1	(-2.4 / +1.7)	13.5	(-1.4 / +1.2)	39.9	(-2.1 / +2.3)
Qwen3-8B + CE-RM-4B (Ours)	77.6	(-1.6 / +1.7)	18.2	(-1.6 / +1.4)	49.1	(-2.1 / +1.9)

Table 5: Evaluation results on Arena-Hard for the original policy models and models trained with different optimizations. The three columns correspond to Arena-Hard v0.1 and the two categories of Arena-Hard v2, evaluating with the official configuration and style control.

often relied on DPO-style optimization, where each preference pair is constructed from a correct and an incorrect evaluation (Mahan et al., 2024; Ye et al., 2025; Yu et al., 2025). More recently, driven by the emergence of reinforcement learning with verifiable rewards (RLVR), research has increasingly shifted toward on-policy optimization (e.g., GRPO), further improving the evaluation performance of GRMs (Chen et al., 2025a; Zhang et al., 2025d; Liu et al., 2025b; Whitehouse et al., 2025).

8 Conclusion

In this work, we start from the limitations of current generative reward models and aim to narrow the gap between their benchmark performance and actual effectiveness in RL practice. To this end, we introduce a two-stage evaluation scheme that enforces unified, query-based criteria across multiple candidate responses. Through a proposed RL method that strengthens the generation of criteria without pointwise label annotations, we train an efficient generative reward model, CE-RM-4B. It features pointwise evaluation and outperforms other much larger pairwise GRMs on common benchmarks and realistic Best-of-N scenarios, while demonstrating better practical RL improvements.

Limitations

Although our proposed method optimizes both criteria and evaluation through a two-stage rollout and achieves superior model performance, there remains room for improving the estimation of the reward signal. For instance, we could use a small

amount of data annotated with reliable pointwise score labels as anchors, thereby providing precise reward signals for the corresponding evaluations during RL to perform calibration. Moreover, since our evaluation protocol takes the form of a multi-turn conversation, and many current LLMs already possess certain tool-use capabilities, it is natural to train the model to leverage tools for more effective evaluation. In fact, we have made some efforts in this direction and found that tool assistance can indeed yield significant gains in evaluation performance for scenarios such as mathematics, code, and factuality; however, it degrades performance on general chat. Nevertheless, these approaches appear promising and can be further explored in future work.

Ethical considerations

This work does not pose any ethical issues. All datasets, open-source LLMs, and API calls used in our work are publicly available. We carefully examine the data used in our work to ensure that it does not contain any personally identifying information or offensive content. And we comply with their respective licenses and use them only for research purposes. In addition, we use an AI assistant to help check the grammatical correctness of this paper.

Acknowledgements

This work was supported by Beijing Natural Science Foundation (L253001). We appreciate the

anonymous reviewers for their helpful comments. Xiaojun Wan is the corresponding author.

References

- Ralph Allan Bradley and Milton E Terry. 1952. Rank analysis of incomplete block designs: I. the method of paired comparisons. *Biometrika*, 39(3/4):324–345.
- Maosong Cao, Alexander Lam, Haodong Duan, Hongwei Liu, Songyang Zhang, and Kai Chen. 2024. Compassjudge-1: All-in-one judge model helps model evaluation and evolution. *arXiv preprint arXiv:2410.16256*.
- Nuo Chen, Zhiyuan Hu, Qingyun Zou, Jiaying Wu, Qian Wang, Bryan Hooi, and Bingsheng He. 2025a. Judgelrm: Large reasoning models as a judge. *arXiv preprint arXiv:2504.00050*.
- Xiuxi Chen, Gaotang Li, Ziqi Wang, Bowen Jin, Cheng Qian, Yu Wang, Hongru Wang, Yu Zhang, Denghui Zhang, Tong Zhang, Hanghang Tong, and Heng Ji. 2025b. RM-R1: reward modeling as reasoning. *CoRR*, abs/2505.02387.
- Arpad E Elo and Sam Sloan. 1978. The rating of chess-players: Past and present. (*No Title*).
- Evan Frick, Tianle Li, Connor Chen, Wei-Lin Chiang, Anastasios Nikolas Angelopoulos, Jiantao Jiao, Banghua Zhu, Joseph E Gonzalez, and Ion Stoica. 2025. How to evaluate reward models for rlhf. In *The Thirteenth International Conference on Learning Representations*.
- Mingqi Gao, Xinyu Hu, Xunjian Yin, Jie Ruan, Xiao Pu, and Xiaojun Wan. 2025. Llm-based nlg evaluation: Current status and challenges. *Computational Linguistics*, pages 1–27.
- Jiaxin Guo, Zewen Chi, Li Dong, Qingxiu Dong, Xun Wu, Shaohan Huang, and Furu Wei. 2025. Reward reasoning model. *arXiv preprint arXiv:2505.14674*.
- Tianxing He, Jingyu Zhang, Tianle Wang, Sachin Kumar, Kyunghyun Cho, James Glass, and Yulia Tsvetkov. 2023. On the blind spots of model-based evaluation metrics for text generation. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 12067–12097.
- Hui Huang, Yancheng He, Hongli Zhou, Rui Zhang, Wei Liu, Weixun Wang, Wenbo Su, Bo Zheng, and Jiaheng Liu. 2025. Think-j: Learning to think for generative llm-as-a-judge. *arXiv preprint arXiv:2505.14268*.
- Aaron Hurst, Adam Lerer, Adam P. Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, Aleksander Madry, Alex Baker-Whitcomb, Alex Beutel, Alex Borzunov, Alex Carney, Alex Chow, Alex Kirillov, Alex Nichol, Alex Paino, and 79 others. 2024. Gpt-4o system card. *CoRR*, abs/2410.21276.
- Nathan Lambert, Jacob Morrison, Valentina Pyatkin, Shengyi Huang, Hamish Ivison, Faeze Brahman, Lester James V. Miranda, Alisa Liu, Nouha Dziri, Shane Lyu, Yuling Gu, Saumya Malik, Victoria Graf, Jena D. Hwang, Jiangjiang Yang, Ronan Le Bras, Oyvind Tafjord, Chris Wilhelm, Luca Soldaini, and 4 others. 2025a. [Tulu 3: Pushing frontiers in open language model post-training](#). *Preprint*, arXiv:2411.15124.
- Nathan Lambert, Valentina Pyatkin, Jacob Morrison, Lester James V. Miranda, Bill Yuchen Lin, Khyathi Raghavi Chandu, Nouha Dziri, Sachin Kumar, Tom Zick, Yejin Choi, Noah A. Smith, and Hananeh Hajishirzi. 2025b. Rewardbench: Evaluating reward models for language modeling. In *NAACL (Findings)*, pages 1755–1797. Association for Computational Linguistics.
- Tianle Li, Wei-Lin Chiang, Evan Frick, Lisa Dunlap, Tianhao Wu, Banghua Zhu, Joseph E Gonzalez, and Ion Stoica. 2024a. From crowdsourced data to high-quality benchmarks: Arena-hard and benchbuilder pipeline. *arXiv preprint arXiv:2406.11939*.
- Tianle Li, Wei-Lin Chiang, Evan Frick, Lisa Dunlap, Banghua Zhu, Joseph E Gonzalez, and Ion Stoica. 2024b. [From live data to high-quality benchmarks: The arena-hard pipeline](#).
- Hunter Lightman, Vineet Kosaraju, Yuri Burda, Harrison Edwards, Bowen Baker, Teddy Lee, Jan Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. 2023. Let’s verify step by step. In *The twelfth international conference on learning representations*.
- Chin-Yew Lin. 2004. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, pages 74–81.
- Chris Yuhao Liu, Liang Zeng, Jiakai Liu, Rui Yan, Jujie He, Chaojie Wang, Shuicheng Yan, Yang Liu, and Yahui Zhou. 2024a. Skywork-reward: Bag of tricks for reward modeling in llms. *arXiv preprint arXiv:2410.18451*.
- Yantao Liu, Zijun Yao, Rui Min, Yixin Cao, Lei Hou, and Juanzi Li. 2025a. Rm-bench: Benchmarking reward models of language models with subtlety and style. In *The Thirteenth International Conference on Learning Representations*.
- Yuxuan Liu, Tianchi Yang, Shaohan Huang, Zihan Zhang, Haizhen Huang, Furu Wei, Weiwei Deng, Feng Sun, and Qi Zhang. 2024b. Calibrating llm-based evaluator. In *Proceedings of the 2024 joint international conference on computational linguistics, language resources and evaluation (Irec-coling 2024)*, pages 2638–2656.
- Zijun Liu, Peiyi Wang, Runxin Xu, Shirong Ma, Chong Ruan, Peng Li, Yang Liu, and Yu Wu. 2025b.

- Inference-time scaling for generalist reward modeling. *arXiv preprint arXiv:2504.02495*.
- Renze Lou, Kai Zhang, and Wenpeng Yin. 2024. Large language model instruction following: A survey of progresses and challenges. *Computational Linguistics*, 50(3):1053–1095.
- Dakota Mahan, Duy Van Phung, Rafael Rafailov, Chase Blagden, Nathan Lile, Louis Castricato, Jan-Philipp Fränken, Chelsea Finn, and Alon Albalak. 2024. Generative reward models. *arXiv preprint arXiv:2410.12832*.
- Saumya Malik, Valentina Pyatkin, Sander Land, Jacob Morrison, Noah A Smith, Hannaneh Hajishirzi, and Nathan Lambert. 2025. Rewardbench 2: Advancing reward model evaluation. *arXiv preprint arXiv:2506.01937*.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, pages 311–318.
- Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea Finn. 2023. Direct preference optimization: Your language model is secretly a reward model. *Advances in neural information processing systems*, 36:53728–53741.
- Ricardo Rei, Craig Stewart, Ana C Farinha, and Alon Lavie. 2020. Comet: A neural framework for mt evaluation. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2685–2702.
- Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, YK Li, Yang Wu, and 1 others. 2024a. Deepseekmath: Pushing the limits of mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300*.
- Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Mingchuan Zhang, Y. K. Li, Y. Wu, and Daya Guo. 2024b. Deepseekmath: Pushing the limits of mathematical reasoning in open language models. *CoRR*, abs/2402.03300.
- Hao Sun, Yunyi Shen, and Jean-Francois Ton. 2025. Rethinking reward modeling in preference-based large language model alignment. In *The Thirteenth International Conference on Learning Representations*.
- Sijun Tan, Siyuan Zhuang, Kyle Montgomery, William Yuan Tang, Alejandro Cuadron, Chenguang Wang, Raluca Popa, and Ion Stoica. 2025. Judgebench: A benchmark for evaluating llm-based judges. In *The Thirteenth International Conference on Learning Representations*.
- Gemini Team. 2025a. Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities. *CoRR*, abs/2507.06261.
- Qwen Team. 2025b. Qwen3-max: Just scale it.
- Qwen Team. 2025c. [Qwen3 technical report](#). *Preprint*, arXiv:2505.09388.
- Jiaan Wang, Yunlong Liang, Fandong Meng, Zengkui Sun, Haoxiang Shi, Zhixu Li, Jinan Xu, Jianfeng Qu, and Jie Zhou. 2023. Is chatgpt a good nlg evaluator? a preliminary study. In *Proceedings of EMNLP Workshop*, page 1.
- Peifeng Wang, Austin Xu, Yilun Zhou, Caiming Xiong, and Shafiq Joty. 2025a. Direct judgement preference optimization. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 1979–2009.
- Tianlu Wang, Ilia Kulikov, Olga Golovneva, Ping Yu, Weizhe Yuan, Jane Dwivedi-Yu, Richard Yuanzhe Pang, Maryam Fazel-Zarandi, Jason Weston, and Xian Li. 2024. Self-taught evaluators. *arXiv preprint arXiv:2408.02666*.
- Weixun Wang, Shaopan Xiong, Gengru Chen, Wei Gao, Sheng Guo, Yancheng He, Ju Huang, Jiaheng Liu, Zhendong Li, Xiaoyang Li, Zichen Liu, Haizhou Zhao, Dakai An, Lunxi Cao, Qiyang Cao, Wanxi Deng, Feilei Du, Yiliang Gu, Jiahe Li, and 22 others. 2025b. Reinforcement learning optimization for large-scale learning: An efficient and user-friendly scaling library. *CoRR*, abs/2506.06122.
- Xueru Wen, Jie Lou, Yaojie Lu, Hongyu Lin, Xingyu, Xinyu Lu, Ben He, Xianpei Han, Debing Zhang, and Le Sun. 2025. Rethinking reward model evaluation: Are we barking up the wrong tree? In *ICLR*. Open-Review.net.
- Chenxi Whitehouse, Tianlu Wang, Ping Yu, Xian Li, Jason Weston, Ilia Kulikov, and Swarnadeep Saha. 2025. J1: Incentivizing thinking in llm-as-a-judge via reinforcement learning. *arXiv preprint arXiv:2505.10320*.
- Ran Xu, Jingjing Chen, Jiayu Ye, Yu Wu, Jun Yan, Carl Yang, and Hongkun Yu. 2025. Incentivizing agentic reasoning in llm judges via tool-integrated reinforcement learning. *arXiv preprint arXiv:2510.23038*.
- Shuying Xu, Junjie Hu, and Ming Jiang. 2024. Large language models are active critics in nlg evaluation. *arXiv preprint arXiv:2410.10724*.
- Ziyi Ye, Xiangsheng Li, Qiuchi Li, Qingyao Ai, Yujia Zhou, Wei Shen, Dong Yan, and Yiqun Liu. 2025. Learning llm-as-a-judge for preference alignment. In *The Thirteenth International Conference on Learning Representations*.
- Jiachen Yu, Shaoning Sun, Xiaohui Hu, Jiayu Yan, Kaidong Yu, and Xuelong Li. 2025. Improve llm-as-a-judge ability as a general ability. *arXiv preprint arXiv:2502.11689*.

- Weizhe Yuan, Graham Neubig, and Pengfei Liu. 2021. Bartscore: Evaluating generated text as text generation. *Advances in neural information processing systems*, 34:27263–27277.
- Lunjun Zhang, Arian Hosseini, Hritik Bansal, Mehran Kazemi, Aviral Kumar, and Rishabh Agarwal. 2025a. Generative verifiers: Reward modeling as next-token prediction. In *The Thirteenth International Conference on Learning Representations*.
- Taolin Zhang, Maosong Cao, Alexander Lam, Songyang Zhang, and Kai Chen. 2025b. Compassjudge-2: Towards generalist judge model via verifiable rewards. *arXiv preprint arXiv:2507.09104*.
- Yanzhao Zhang, Mingxin Li, Dingkun Long, Xin Zhang, Huan Lin, Baosong Yang, Pengjun Xie, An Yang, Dayiheng Liu, Junyang Lin, Fei Huang, and Jingren Zhou. 2025c. Qwen3 embedding: Advancing text embedding and reranking through foundation models. *arXiv preprint arXiv:2506.05176*.
- Yifan Zhang, Xingyu Lu, Xiao Hu, Chaoyou Fu, Bin Wen, Tianke Zhang, Changyi Liu, Kaiyu Jiang, Kaibing Chen, Kaiyu Tang, Haojie Ding, Jiankang Chen, Fan Yang, Zhang Zhang, Tingting Gao, and Liang Wang. 2025d. R1-reward: Training multimodal reward model through stable reinforcement learning. *CoRR*, abs/2505.02835.

A Evaluation Prompts

We provide our evaluation prompts in Figures 4–6.

The prompt for the first evaluation setting and Eq. (1) in Section 2

Your task is to evaluate the quality of a response to the given user query.

Provide your evaluation with a careful and comprehensive analysis, followed by a corresponding overall quality score from 0 to 10 within .

Use integers or half-point increments for the score, with higher numbers representing higher quality.

Below are the user query and the response:

[Start of Query]
{instruction}
[End of Query]

[Start of Response]
{response}
[End of Response]

Figure 4: The prompt for the first evaluation setting and Eq. (1) in Section 2.

The prompt for the second evaluation setting and Eq. (2) in Section 2

Your task is to evaluate the quality of a response to the given user query.

Begin by carefully analyzing the query to fully understand the user's intent and requirements, and then take into account all common and tangible factors that can indicate the response quality.

From these considerations and analyses, derive the final evaluation criteria list between [Start of Criteria] and [End of Criteria], with one criterion per line.

Next, for each criterion, focus on its concerns and carefully evaluate the corresponding specific quality of the response, providing the detailed analysis as well as relevant arguments, followed by the corresponding quality score from 0 to 5 within .

Finally, based on the analyses of these criteria, including their relative importance and scores, conduct a comprehensive evaluation of the response's overall quality with sufficient and explicit evidence, and then provide a corresponding overall quality score from 0 to 10 within .

Use integers or half-point increments for all scores, with higher numbers representing higher quality.

Below are the user query and the response:

[Start of Query]
{instruction}
[End of Query]

[Start of Response]
{response}
[End of Response]

Figure 5: The prompt for the second evaluation setting and Eq. (2) in Section 2.

The prompts for the third evaluation setting and Eq. (3) in Section 2

Stage 1

Your task is to produce a minimal set of criteria for evaluating the quality of potential responses to the user query given below.

Begin by carefully analyzing the query to fully understand the user's intent and requirements, and then take into account all common and tangible factors that can indicate the response quality.

From these considerations, derive the final evaluation criteria list, which must adhere to the following requirements:

- Each criterion should consist of a concise term as well as its unambiguous description.
- The number of criteria is not necessarily the more the better; Fewer yet comprehensive is more desired.
- The criteria should be sufficient and complete, ensuring that no essential aspects or key signals of response quality are omitted.
- The criteria should be necessary and non-overlapping, with each one indispensable, distinct in perspective, and strictly orthogonal to others.

Provide the relevant analysis first, followed by the numbered list of criteria between [Start of Criteria] and [End of Criteria], with one criterion per line and the more important ones coming first.

Below is the user query:

[Start of Query]
{instruction}
[End of Query]

Stage 2

Now that you have a response to the previous user query, your new task is to evaluate it using the criteria list you have produced.

For each criterion, focus on its concerns and carefully evaluate the corresponding specific quality of the response, providing the detailed analysis as well as relevant arguments, followed by the corresponding quality score from 0 to 5 within .

Moreover, if the response demonstrates strengths or weaknesses beyond the scope of your criteria list, introduce an additional criterion titled "Other Point(s)," discussing them and considering them as bonus points or deductions as appropriate.

Finally, based on the analyses of these criteria, including their relative importance and scores, conduct a comprehensive evaluation of the response's overall quality with sufficient and explicit evidence, and then provide a corresponding overall quality score from 0 to 10 within .

Use integers or half-point increments for all scores, with higher numbers representing higher quality.

Below is the response:

[Start of Response]
response
[End of Response]

Figure 6: The prompts for the third evaluation setting and Eq. (3) in Section 2.