

# Reward Modeling for Scientific Writing Evaluation

Furkan Şahinuç<sup>1,2</sup>, Subhabrata Dutta<sup>1</sup>, Iryna Gurevych<sup>1,2</sup>

<sup>1</sup>Ubiquitous Knowledge Processing Lab (UKP Lab)

Department of Computer Science and Hessian Center for AI (hessian.AI)

Technical University of Darmstadt

<sup>2</sup>Konrad Zuse School of Excellence in Learning and Intelligent Systems (ELIZA)

[www.ukp.tu-darmstadt.de](http://www.ukp.tu-darmstadt.de)

## Abstract

Scientific writing is an expert-domain task that demands deep domain knowledge, task-specific requirements and reasoning capabilities that leverage the domain knowledge to satisfy the task specifications. While scientific text generation has been widely studied, its evaluation remains a challenging and open problem. It is critical to develop models that can be reliably deployed for evaluating diverse open-ended scientific writing tasks while adhering to their distinct requirements. However, existing LLM-based judges and reward models are primarily optimized for general-purpose benchmarks with fixed scoring rubrics and evaluation criteria. Consequently, they often fail to reason over sparse knowledge of scientific domains when interpreting task-dependent and multi-faceted criteria. Moreover, fine-tuning for each individual task is costly and impractical for low-resource settings. To bridge these gaps, we propose cost-efficient, open-source reward models tailored for scientific writing evaluation. We introduce a two-stage training framework that initially optimizes scientific evaluation preferences and then refines reasoning capabilities. Our multi-aspect evaluation design and joint training across diverse tasks enable fine-grained assessment and robustness to dynamic criteria and scoring rubrics. Experimental analysis shows that our training regime strongly improves LLM-based scientific writing evaluation. Our models generalize effectively across tasks and to previously unseen scientific writing evaluation settings, allowing a single trained evaluator to be reused without task-specific retraining. We make our code<sup>1</sup> and data<sup>2</sup> publicly available.

## 1 Introduction

Due to the strong text generation capabilities of large language models (LLMs), their application

<sup>1</sup>GitHub: [UKPLab/acl2026-expert-rm](https://github.com/UKPLab/acl2026-expert-rm)

<sup>2</sup>Data: [TUdatalib](https://TUdatalib)

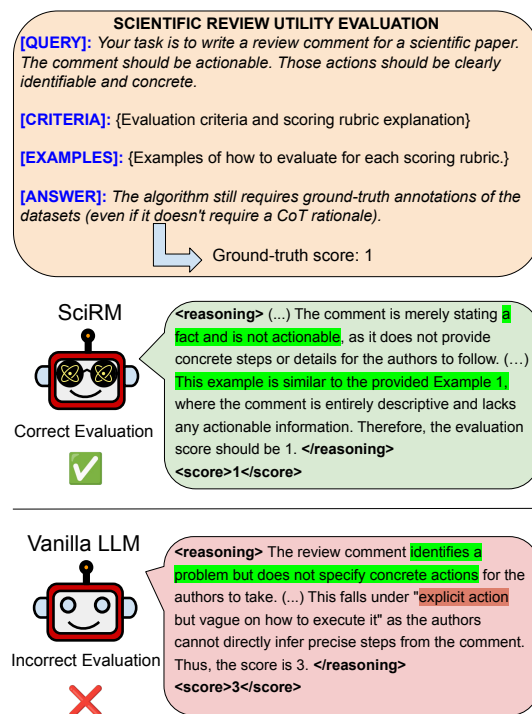


Figure 1: Example demonstration of how we formalize scientific writing evaluation and the outputs of the review utility evaluation task. Vanilla LLM-based judges fail to properly reason over the task-specific evaluation criteria and provided examples. Contradictory statements are highlighted in different colors. In contrast, our SciRM model successfully incorporates the given criteria and examples into its reasoning process and correctly evaluates the scientific artifact.

to scientific text generation, such as related work generation, review generation, and revising papers, has gained increasing attention recently (Li and Ouyang, 2024; Liang et al., 2024; Afzal et al., 2026). However, without proper evaluation, it is difficult to assess the accuracy and reliability of the generated texts. Therefore, the limitations in evaluation can block the entire development pipeline. Since scientific writing tasks have diverse and task-specific requirements, developing appropriate eval-

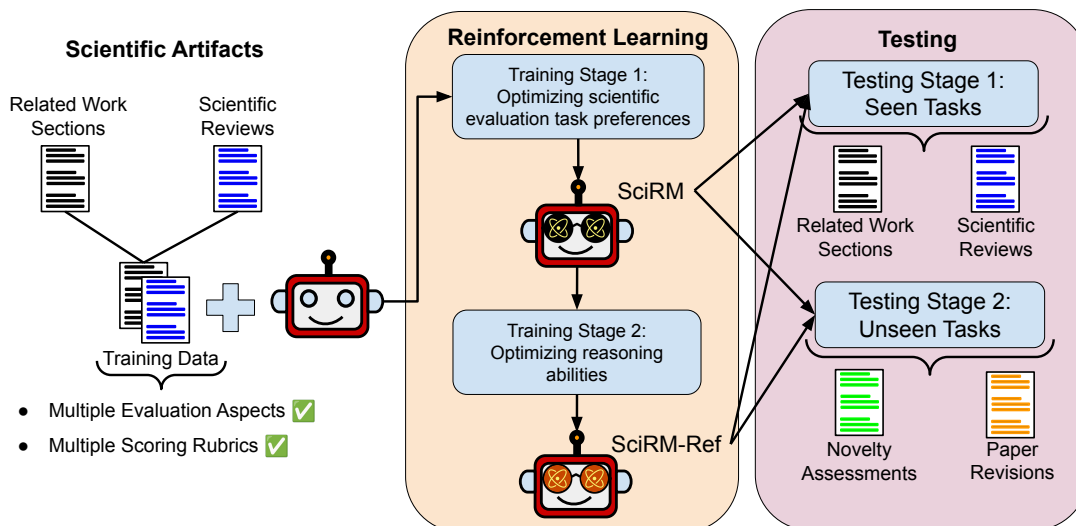


Figure 2: Overview of SCIRM and SCIRM-REF training and testing pipeline. Diverse scientific artifacts are used to construct training data with multiple evaluation aspects and scoring rubrics (see Section 3.1 for details). Models are trained via GRPO in two stages to optimize task specifications and reasoning capabilities, and are evaluated on both seen and unseen scientific writing evaluation tasks.

uation frameworks is a challenging problem. Training evaluators for each individual task is costly and, in some cases, infeasible due to limited data availability.

LLM-as-a-judge approaches (Liu et al., 2023; Zheng et al., 2023) are the most widely adopted evaluation paradigms in scientific writing tasks. However, they often fail to reason over the given domain knowledge and task-specific preferences (see Figure 1). This motivates the need for mechanisms that allow models to reason over and remain grounded in explicit evaluation guidelines (aka, *constitution*) at inference time. Inference-time adaptability is a major challenge for existing approaches like Constitutional AI (Bai et al., 2022), which internalize fixed constitution during training and therefore cannot be readily applied to a diverse set of evaluation tasks. This rigidity is problematic while evaluating scientific text generations, as evaluation guidelines can be different, even contradictory, across different aspects, tasks, or domains.

To improve the reasoning capabilities of LLM-based judges, training reward reasoning models has recently gained popularity (Ankner et al., 2024; Guo et al., 2025b; Chen et al., 2025, 2026; Wang et al., 2024). However, existing reward models are primarily optimized for improving performance on community-standard benchmarks, such as mathematical reasoning, instruction-following, and human-preference modeling for coding, helpfulness, and safety (Lambert et al., 2025; Malik

et al., 2026; Frick et al., 2025). Therefore, they fall short of capturing the nuanced requirements of scientific writing evaluation. In addition, the vast majority of reward models encode task preferences in a pairwise manner, which prevents independent assessment of text quality based on explicit task-specific criteria. Another drawback of such models is that they are optimized for fixed scoring rubrics and criteria. Scientific tasks have unique characteristics; each requires domain-specific expertise and evaluation dynamics that differ from those of open-ended creative writing tasks (Chakrabarty et al., 2025). Even for the same scientific artifact, it is possible to evaluate from multiple aspects requiring different criteria and rubrics. However, current reward models experience performance degradation when they are applied to other tasks with evaluation rubrics that differ from those seen during training (Yang et al., 2024). On the other hand, Şahinuç et al. (2025) offers insights on multiple-aspect evaluation for expert-domain tasks. However, their approach is limited to a single task, and their most accurate pipeline relies on proprietary LLMs, which restricts scalability and hinders generalization across diverse scientific writing tasks.

In this work, we adapt reward model training strategies to enhance LLMs’ capabilities regarding *what to evaluate* and *how to evaluate* scientific writing tasks. Concretely, we design reward models that are conditioned on an explicit evaluation constitution—a structured description of criteria and

label space—present during both training and inference. Furthermore, we introduce a two-stage optimization process in which models not only learn to follow the constitution but also reflexively reinterpret it to correct and stabilize their own reasoning. This process involves a joint optimization of in-context preference-following and reasoning abilities, which (Lai et al., 2024) identifies as a missing piece in modern LMs trained using RL.

**Contributions and Findings:** We introduce cost-efficient reward models, SCIRM and SCIRM-REF, specifically designed for scientific writing evaluation (C1). We employ two-stage reinforcement learning to optimize the models for (1) scientific writing evaluation preferences and (2) reasoning abilities to better comprehend the given evaluation criteria, enabling models to explicitly reason over and faithfully adhere to dynamically specified evaluation rules. (C2). Rather than producing a single aggregated score, our models evaluate scientific artifacts across multiple aspects. This approach enhances both reliability and interpretability of our evaluation (C3). We curate and process datasets from diverse sources and jointly train our models across multiple tasks to (1) improve robustness against varying scoring rubrics and (2) enhance our models’ generalization capabilities (C4). We illustrate the overview of our pipeline in Figure 2.

We finally test our models on four different scientific writing tasks: related work sections, paper reviews, novelty summary assessments, and paper revisions based on instructions, each with distinct evaluation aspects and scoring rubrics. Experimental results show that our two-stage training scheme substantially boosts the LLMs’ scientific writing evaluation performance (F1). In particular, our second training stage leads to improvements in tasks requiring strong reasoning capabilities (F2). Furthermore, our models outperform baseline models on tasks that are not included in the training, indicating our model’s strong generalization and scaling capabilities (F3).

## 2 Related Work

### 2.1 Scientific Writing Evaluation

Utilizing LLMs as direct evaluators is one of the most intuitive approaches for scientific writing evaluation due to their flexibility in being prompted with various task settings without additional training overhead (Liu et al., 2023; Zheng et al., 2023). However, prior work has shown that vanilla LLM-

as-a-judge setups are prone to systematic biases and failures in domain-grounded reasoning (Li et al., 2024; Szymanski et al., 2025; Gao et al., 2025). Moreover, the relative scarcity of scientific writing evaluation datasets also prevents LLMs from becoming more familiar with these tasks in the training phase.

To address these limitations, Jourdan et al. (2025) focus on evaluating scientific text revision. They highlight that LLM-as-a-judge methods struggle to grasp task-specific evaluation aspects in the absence of a gold reference. In a complementary direction, Purkayastha et al. (2025) classify the peer-reviews according to common mistakes of reviewers referred to as *lazy-thinking* patterns. A key limitation of their classification scheme is that, although some review sentences are suitable for multiple error categories, the dataset imposes a one-to-one label assignment. Similarly, Sadallah et al. (2025) introduce a direct-evaluation benchmark that measures the utility of scientific reviews across the aspects of actionability, grounding, verifiability, and helpfulness.

Addressing a different task, Şahinuç et al. (2025) propose a fine-grained evaluation framework for related work generation. Instead of producing a single overall score, they perform aspect based evaluation. Although their evaluation achieves strong alignment with human experts, their implementation is limited to the related work generation task.

### 2.2 Evaluation-Tuned Models

Reward modeling has a significant impact on the success of post-training with reinforcement learning. Although verifiable rewards are computationally efficient and demonstrate strong performance in math and coding tasks (Wen et al., 2026), many complex tasks are not suitable for directly verifiable reward signals, such as writing quality or helpfulness. These limitations and the recent success of reasoning-centric models on complex tasks (Guo et al., 2025a) have motivated the development of reward reasoning models (Ankner et al., 2024; Guo et al., 2025b; Chen et al., 2025, 2026; Wang et al., 2024). The primary objective of such models is to generate more reliable and accurate rewards in non-verifiable tasks by leveraging intermediate thinking steps. Along with reasoning features, prior works also explore improving long-context comprehension (Tang et al., 2025), integrating external documents (Ma et al., 2025), and bridging pointwise and pairwise scoring paradigms in reward modeling set-

tings (Whitehouse et al., 2026; Jian et al., 2025). The main drawback of these models is that they are optimized for standard reward benchmarks, which do not involve any scientific writing tasks. Therefore, they struggle to adapt to the varying scientific writing criteria and scoring.

Although there are attempts to automatically generate task-specific evaluation criteria alongside the evaluation itself (Liu et al., 2025a,b; Liang et al., 2025), the generated criteria are mostly surface-level and do not match the scientific writing evaluation requirements which are highly specific, diverse, and involve domain expertise.

Besides reward models, evaluation-tuned LLM-as-a-judge models also exist. In general, these models are not specialized for a particular task but are designed as general-purpose judge models (Liu et al., 2023; Kim et al., 2024; Alexandru et al., 2025; Flow AI, 2024; Shiwen et al., 2024). This feature provides flexibility, allowing them to be employed in a diverse set of tasks. However, since they are not specialized for expert-domain tasks like scientific writing, they experience performance drops. In contrast, our work directly focuses on the evaluation of scientific writing and generalizes across different tasks with unique criteria and evaluation schemas.

### 3 Methodology

#### 3.1 Dataset

To improve the generalization capabilities of our reward models, we include various scientific writing tasks that can be evaluated from multiple aspects using different scoring rubrics. These tasks focus on related work and review generation. Şahinuç et al. (2025) provide a fine-grained set of aspects and principles for evaluating LLM-generated related work sections. From these aspects, we include coherence (whether citation sentences are consistent with the cited papers), positioning type (whether related works position themselves within existing literature, given a specific type), and positioning consistency (how well this positioning is maintained across different paragraphs) in our dataset. These aspects are subject to binary evaluation. To increase data diversity, we expand their dataset (originally containing NLP papers) by incorporating papers from other top-tier machine learning venues. We follow their preprocessing methodology to construct individual data instances.

Additionally, we incorporate the RevUtil dataset

(Sadallah et al., 2025). The tasks in this dataset are evaluated on a 1–5 scale, with higher scores reflecting better quality. The number of instances and the corresponding scoring rubrics for each aspect of the training tasks are given in Appendix A. In addition, example data instances can be found in evaluation prompts in Appendix C.

#### 3.2 Reward Model Training

Our training scheme consists of two distinct sequential stages: (1) teaching model to evaluate a candidate scientific text by reasoning upon an in-context constitution and (2) intensifying the reasoning acquired at stage 1 by teaching model self-verification based on the constitution. In both stages, we employ the GRPO algorithm due to its effectiveness in preference optimization and reasoning enhancement (Shao et al., 2024). We formulate our task as follows: Given a prompt  $x = (q, c, e)$ , consisting of a task query  $q$ , criteria (or constitution) with a scoring rubric  $c$  and example evaluations for each possible score  $e$ , generate an evaluation output  $M(x) = y = (j, s)$  that consists of a reasoning (or justification) component  $j$  along with a final score (or verdict)  $s$ .

To facilitate trivial extraction of rewards from the model outputs across different tasks and scoring rubrics, we use formatted outputs where the reasoning and scoring parts are encapsulated between `<reasoning>` `</reasoning>` and `<score>` `</score>` tags, respectively. The reasoning part helps to improve evaluation performance, as shown in previous works (Kim et al., 2025; Whitehouse et al., 2026; Ma et al., 2025). We calculate our rewards based on compliance with the formatting and correctness of the generated score. In particular, let  $y$  be the output, and let  $m(y)$  be the score assigned by the model within `<score>` `</score>` in  $y$ , or  $m(y) = \emptyset$  if no valid tag is present. Let  $s^*$  be the gold score and let  $S$  be the set of possible scores in the corresponding scoring rubric. The reward function for the first training stage is defined as:

$$r(y, s^*, S) = \begin{cases} -0.5 & \text{if } m(y) = \emptyset \\ 0 & \text{if } m(y) \notin \mathbb{R} \\ 0.25 & \text{if } m(y) \notin S \\ 0.5 & \text{if } m(y) \in S \ \& \ m(y) \neq s^* \\ 1.5 & \text{if } m(y) \in S \ \& \ m(y) = s^* \end{cases}$$

This rewarding strategy helps identify different error types, including formatting and semantic errors. Our initial experiments showed that models

may attempt to exploit the reward function (Skalse et al., 2022) by generating responses that only focus on the scoring criteria. We also observed occasional instability during training, where models would fail to finish their outputs and continue generating until the maximum length was reached. To mitigate these issues, we introduce a length penalty function to mitigate training instability and reward hacking.

$$f(L, T) = \begin{cases} \frac{(kT-L)^2}{kT^2}, & L < kT, \\ \frac{(L-(1-k)T)^2}{(T-(1-k)T)^2}, & L > (1-k)T, \\ 0, & \text{otherwise} \end{cases}$$

Here,  $L$  and  $T$  stand for the number of generated tokens and the maximum token length, respectively. The parameter  $k \in (0, 0.5]$  controls how strict the penalty will be. When the generated output is too short or too long—i.e., closer to the boundary regions, the penalty quadratically increases with the distance from the acceptable range. If the length is within this range, no penalty is applied. The final reward is calculated by

$$R(x) = r(M(x), s^*, S) - f(L, T).$$

To further enhance the reasoning capabilities of our model, we introduce a second reinforcement learning stage, where the model is incentivized to update its evaluation by reflecting on the constitution. We first obtain the scores and reasoning from the model on the data instances. Next, we remove the scoring part and append the reasoning part with a reflection prompt that encourages the model to revisit the query and criteria before outputting its final score. The adjusted reward function for the second stage training is as follows:

$$r_p(s_i, s_f, s^*) = \begin{cases} 0.5 & \text{if } s_i = s^* \ \& \ s_f = s^* \\ 1.0 & \text{if } s_i \neq s^* \ \& \ s_f = s^* \\ -0.5 & \text{if } s_i \neq s^* \ \& \ s_f \neq s^* \\ -1.0 & \text{if } s_i = s^* \ \& \ s_f \neq s^* \end{cases}$$

$s_i$  and  $s_f$  denote the initial and final scores, respectively. This rewarding strategy depends not only on the correctness of the final decision but also on the model’s prior performance. If the model corrects itself, it receives the highest reward; when it backslides, it is harshly penalized. Consistent behaviors yield a moderate reward or penalty based

on correctness. In other words, we promote the model to revisit the criteria and its earlier reasoning either to fix the incorrect trajectory or to maintain the correct one. We refer to the model trained only in the first stage as SCIRM, and the model trained using the two-stage procedure as SCIRM-REF.

## 4 Experiments

**Baselines:** We use Qwen 2.5-7B (Yang et al., 2025b) as our base model, based on its performance in our preliminary experiments on scientific writing evaluation. We evaluate our model against several open-weight LLMs of comparable size, representing different model families: DS-R1-Qwen-7B (Guo et al., 2025a), Granite3.3-8B-Instruct (IBM, 2025), Llama3.1-8B-Instruct (Grattafiori et al., 2024), Qwen2.5-7B-Instruct (Yang et al., 2025b), and Qwen3-8B (Yang et al., 2025a). In addition to open-weight models, we experiment with GPT-5.2 (2025-12-11) and o3-mini (2025-01-31), which was used as a backbone model in related work evaluation by Şahinuç et al. (2025), to gain insight about the performance bounds of large proprietary models.

We also include specialized evaluation models in our baseline set. To represent tuned, general purpose LLM-as-a-judge models, we use Prometheus<sup>3</sup> (Kim et al., 2024) and Selene<sup>4</sup> (Alexandru et al., 2025) frameworks. Furthermore, based on open availability and compatibility with pointwise scoring, we include DeepSeek-GRM<sup>5</sup> (Liu et al., 2025b) and Skywork-Critic<sup>6</sup> (Shiwen et al., 2024) as reward modeling baselines.

**Experimental Setting:** We employ LoRA framework (Hu et al., 2022) to meet the memory requirements and facilitate training. Besides testing our reward models on evaluation of related works and scientific reviews, we introduce two additional experimental settings. For each experiment, the inference is repeated five times (GPT-5.2 and o3-mini three times) on test sets with a temperature of 1 and a top-p of 0.95, and we report the resulting mean and standard deviation. Further experimental details and evaluation prompts are provided in Appendix B and C respectively.

<sup>3</sup>M-Prometheus-7B

<sup>4</sup>Selene-1-Mini-Llama-3.1-8B

<sup>5</sup>DeepSeek-GRM-16B

<sup>6</sup>Skywork-Critic-Llama-3.1-70B

## 4.1 Unseen Aspect Experiments

To evaluate whether our training scheme provides an improvement over the baselines on unseen aspects of the same evaluation task, we conduct an experiment that excludes selected evaluation aspects from the training data. Specifically, we remove the *Actionability* and *Grounding* aspects from the RevUtil dataset and test whether models trained on the remaining aspects can still improve performance on unseen evaluation aspects. We refer to the resulting model as SCIRM-MASKED.

## 4.2 Unseen Task Experiments

To assess the generalization capabilities of our reward model, we test it on two additional tasks that are not included in the training data: comparison of novelty assessments and evaluation scientific revision quality. For the first task, we utilize the novelty-evaluation dataset from (Afzal et al., 2026), which contains summaries of both human and LLM-generated novelty assessments. Human summaries are accompanied by human-annotated binary labels indicating the novelty of the papers. However, LLM-generated assessment summaries do not include final novelty labels. To address this, we use GPT-5.1 (OpenAI, 2025) to infer the novelty labels from the LLM-generated summaries. Our evaluation objective is to assess whether LLM-generated summaries align with human summaries in terms of their novelty verdicts (i.e., *novel*, *not novel*). We treat equality of the novelty labels as the ground-truth indicator of alignment between summary pairs. In the final evaluation dataset, there are 76 assessment pairs with 56 aligned pairs and 20 non-aligned pairs.

For the second task, we use the scientific revision dataset from (Jourdan et al., 2025), which includes paragraphs from scientific papers and corresponding LLM-generated revisions for a given instruction. Each candidate revision is annotated along two dimensions: whether it addresses the given instruction (i.e., *relatedness*) and whether it correctly applies the instruction to produce an improved version of the original text (i.e., *correctness*). For each aspect, we have 3,092 data instances, resulting in a total of 6,184 data points.

## 5 Results

**Performance on seen domains.** We begin by evaluating our models on the tasks included in the training phase. For review utility evaluation, we present

results in Figure 3. We observe that our training scheme leads to a clear improvement over baseline models. Both SCIRM and SCIRM-REF consistently achieve the highest accuracy across all evaluation aspects. Among these aspects, all models reach their highest scores on *Verifiability Extraction*. Since this aspect focuses on binary claim detection in reviews rather than scoring on a 1–5 scale, it constitutes a comparatively simpler evaluation than the other aspects, which require a deeper understanding of the evaluation criteria. Across the remaining aspects, our models perform comparably among themselves, which is also observed for the baseline models.

We report the model performances for the related work evaluation in Figure 4. We observe a similar pattern to the results in the review utility evaluation in terms of improvements over the baseline models. In particular, our models achieve near-perfect accuracy on *Positioning Consistency* aspect. In contrast, *Coherence* stands out as a more challenging aspect, since comparing citation sentences with the cited papers requires a higher degree of semantic understanding. Although the performance gap between our models and the baselines is largest for *Positioning Type*, Qwen3, GPT-5.2, and o3-mini exhibit outlier behavior on this aspect, possibly due to better leveraging their reasoning capabilities. However, our models still outperform all baselines when all data instances are considered.

Figures 3 and 4 collectively indicate that our training regime helps to leverage the comprehension capabilities of LLM models in scientific writing evaluation. Furthermore, our joint training across multiple tasks improves the models’ adaptability to diverse scoring rubrics and evaluation criteria. For completeness, we report full numerical results, including mean and standard deviation values for accuracy and F1 score in Appendix D.

**Unseen Aspect Performance.** We further evaluate performance on unseen aspects by reporting the results of SCIRM-MASKED in Figure 3. Since the *Actionability* and *Grounding* aspects are excluded from the training, the performance on these aspects decreases compared to our other models. Nevertheless, SCIRM-MASKED continues to outperform the baseline models, with the exception of Qwen3, GPT-5.2, and o3-mini on the *Grounding* aspect. This behavior suggests that our methodology enables the model to capture the underlying structure of the task, rather than overfitting to specific evaluation aspects observed during training. As a result,

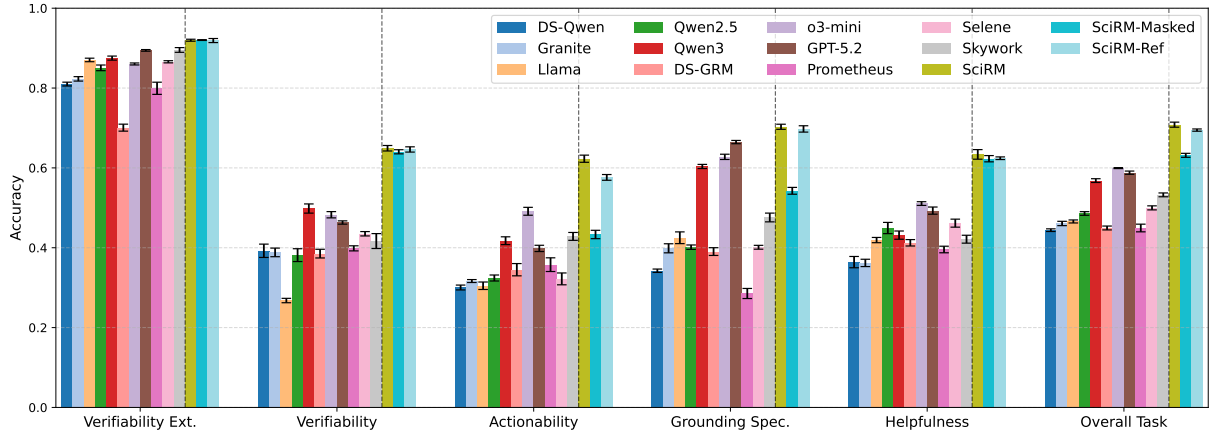


Figure 3: Review utility evaluation results. Results to the left of the vertical dashed lines correspond to the baseline models, whereas results to the right correspond to our models.

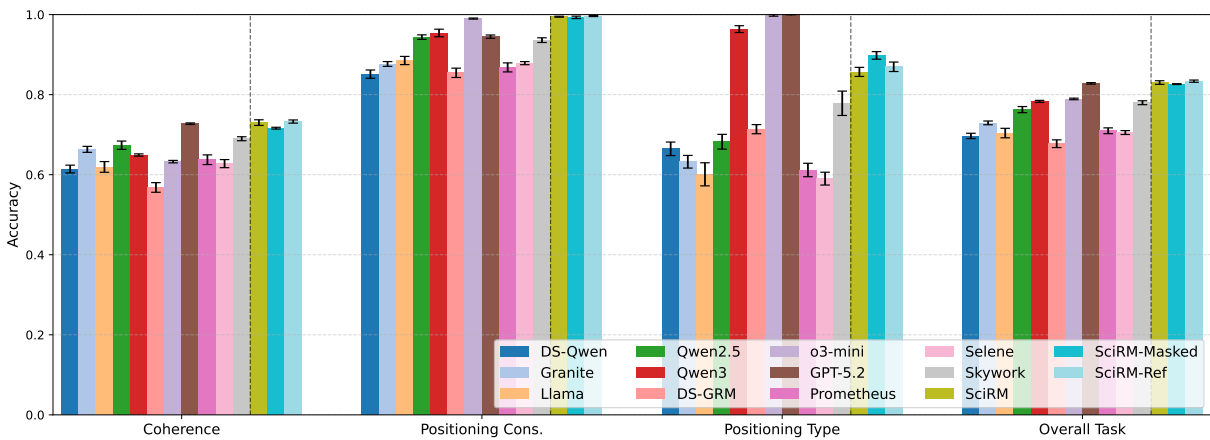


Figure 4: Related work evaluation results. Results to the left of the vertical dashed lines correspond to the baseline models, whereas results to the right correspond to our models.

the model can generalize the learned reasoning patterns and adapt effectively to previously unseen evaluation criteria, even when explicit supervision for those aspects is absent. For the other aspects included in training for both revision utility and related work, SCIRM-MASKED performs similar to our other model variants.

**Unseen Task Performance.** Table 1 shows the performance of our models on tasks that are used exclusively as test sets. In the novelty alignment evaluation, SCIRM-REF not only outperforms the baseline models but also shows a clear improvement over SCIRM. Among all models, o3-mini achieves the highest accuracy, while SCIRM-REF delivers near-closed-model performance despite having smaller size and being fully open. We note that novelty evaluation differs from other tasks, as determining the novelty of a paper based on summarized reviewer comments and comparing these according to their implicit novelty judgments re-

Model	Novelty Align.	Revision		
		Rel.	Cor.	Overall
DS-Qwen	0.41	0.77	0.67	0.72
Granite3.3	0.48	0.80	0.68	0.74
Llama3.1	0.56	0.56	0.61	0.58
Qwen2.5	0.71	0.80	0.66	0.73
Qwen3	0.58	0.86	0.76	0.81
Prometheus	0.33	0.76	0.70	0.73
Selene	0.57	0.79	0.71	0.75
Skywork	0.42	0.83	0.77	0.80
DS-GRM	0.35	0.73	0.59	0.66
o3-mini	0.78	0.88	0.80	0.84
GPT-5.2	0.74	0.82	0.78	0.80
SciRM	0.61	0.82	0.73	0.78
SciRM-Ref	0.74	0.82	0.73	0.78

Table 1: Model performance in terms of accuracy for unseen tasks. Rel. and Cor. denotes Relatedness and Correctness aspects for revision evaluation task.

Model	RW	RU	Nov.	Rev.	Avg.
DS-Qwen	0.70	0.44	0.41	0.72	0.57
Granite3.3	0.73	0.46	0.48	0.74	0.60
Llama3.1	0.70	0.47	0.56	0.58	0.58
Qwen2.5	0.76	0.49	0.71	0.73	0.67
Qwen3	0.78	0.57	0.58	0.81	0.68
Prometheus	0.71	0.45	0.33	0.73	0.56
Selene	0.71	0.50	0.57	0.75	0.63
Skywork	0.78	0.53	0.42	0.80	0.63
DS-GRM	0.68	0.45	0.35	0.66	0.54
o3-mini	0.79	0.60	0.78	0.84	0.75
GPT-5.2	0.83	0.59	0.74	0.80	0.74
SciRM	0.83	0.71	0.61	0.78	0.73
SciRM-Ref	0.83	0.69	0.74	0.78	0.76

Table 2: Individual and average task performances. (RW: Related work, RU: Review utility, Nov: Novelty alignment, Rev: Revision alignment, Avg: Average)

quires strong reasoning capabilities. These results indicate that the second training stage is effective in further enhancing the model’s reasoning ability beyond what is achieved in the initial training stage. Although we provide explicit criteria to illustrate how the evaluation should be performed, only SCIRM-REF is able to effectively leverage this additional supervision.

In the paper revision evaluation, our models either outperform baselines or perform comparably to the proprietary and larger models. We observe that models generally tend to perform better on the *Relatedness* aspect than on *Correctness*. We note that this task primarily evaluates instruction-following on scientific text. Since it involves instructions like *content addition or removal*, *summarization*, and *merging*, implementing evaluations is relatively less reasoning-intensive than in novelty alignment.

Overall, the observed performance gains indicate improved cross-task generalization rather than memorization. Since the unseen tasks are diverse and substantially different from those in training, consistent improvements under this distribution shift suggest that the model has learned transferable capabilities instead of relying on task-specific patterns. These capabilities enable the model to apply learned evaluation principles across diverse scientific writing scenarios.

In Table 2, we summarize the results for all four evaluation tasks. Our reward models outperform the baselines on individual tasks and result in a

Models	Reasoning Difference	Example
SciRM Qwen2.5	Grounded reasoning vs speculative inference	Table 14
SciRM SciRM-Ref	Rubric-alignment vs semantic misweighting	Table 15
SciRM-Ref DS-GRM	Detail sensitivity vs content omission	Table 16

Table 3: Reasoning trace comparison summary table. Corresponding full trace tables are provided in Example column.

larger average score. SCIRM-REF becomes the prominent model by achieving the largest average score. Although it exhibits very similar results to SCIRM on the related work, review utility, and revision evaluation tasks, the difference arises from the novelty alignment evaluation task, in which the reasoning capabilities of the models are crucial.

## 6 Model Analysis and Ablation Studies

### 6.1 Reasoning Traces and Error Analysis.

To obtain further insight into the mechanism behind the improvements of our models, we make a comprehensive analysis on generated reasoning traces across various models and tasks. First, we observe that baseline models tend to make assumptions beyond the provided evaluation criteria rather than fully grounding their judgments in them, as our model does (e.g., inferring implicit actionability despite the absence of explicit actionable guidance.).

Secondly, our second-stage training improves rubric alignment in evaluation, particularly for the novelty alignment task. Without reasoning-focused training, models may overemphasize certain aspects of the assessments (e.g., their characterization as incremental contributions) while disregarding other explicit novelty statements, leading to incorrect conclusions about alignment.

Finally, baseline models can make more direct errors, such as failing to fully comprehend the evaluated text, for example by overlooking the changed text in the scientific revision task. Furthermore, our manual analysis shows that the reasoning traces generated by our models are fully consistent with their final verdicts. Table 3 summarizes the qualitative differences in generated reasoning and provides references to the full reasoning traces in Appendix E.

Model	RW	RU	Nov.	Rev.	Avg.
SciRM	-0.03	-0.06	-0.21	-0.00	-0.07
SciRM-Ref	-0.02	-0.05	-0.25	-0.01	-0.08

Table 4: Ablation study on the reasoning process across tasks. Numbers denote the performance change relative to the full pipeline (Table 2) when reasoning is disabled at evaluation time.

## 6.2 Ablation Study

To investigate the role of the reasoning to guide the final judgment, we ablate the reasoning process for our proposed models. Specifically, we remove the instructions for reasoning process in the system prompt and example reasoning parts in evaluation constitution. In Table 4, we observe that both individual and overall average task performance decrease in the absence of reasoning traces. This suggests that reasoning-conditioned final judgments lead to more accurate results. Importantly, novelty alignment evaluation suffers most from the ablation for both models, as it requires non-trivial generalization of reasoning and grounding into the constitution. This further strengthens our claim that our method introduces generalization to unseen tasks.

## 7 Conclusion

In this paper, we address the problem of evaluating scientific writing tasks. Owing to their unique characteristics, both the generation and evaluation of scientific writing require greater attention than common, well-defined NLP tasks. Moreover, this uniqueness often results in limited annotated data, making it impractical to build specialized models for individual tasks. To address this challenge, we propose reward models that enhance reasoning capabilities and capture task dependencies through explicit evaluation criteria and scoring rubrics. Our models outperform strong baselines in both seen and unseen task settings, highlighting their reasoning and generalization capabilities and underscoring their practicality for real-world scientific workflows. This design enables a single trained evaluator to generalize across tasks, substantially reducing the need for task-specific retraining.

We believe that this work takes a step toward more reliable, interpretable, and reusable evaluation frameworks for scientific writing. A promising direction for future work is to improve scientific text generation across diverse tasks using reinforcement learning, where rewards are provided by our

models. By producing fine-grained, aspect-level feedback on scientific text quality, our models can guide generators along multiple evaluation dimensions rather than optimizing a single objective. Finally, the cost efficiency of our open-source models further enhances their practicality compared to proprietary LLM-based evaluators.

## Limitations

Due to the computational requirements of GRPO, we restrict our experiments to 7B-scale models. We hypothesize that larger models could better capture task-specific evaluation nuances, generalize more effectively, and exhibit more accurate reasoning. Nevertheless, our training scheme provides an opportunity to develop small-sized evaluation models that can outperform or remain competitive with proprietary models. Another limitation is the scarcity of evaluation datasets for scientific writing. This results in most of the scoring rubrics being binary. With the introduction of more fine-grained scientific writing evaluation datasets, a more comprehensive analysis can be conducted as future work.

## Acknowledgments

This work has been funded/co-funded by the European Union (ERC, InterText, 101054961). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council. Neither the European Union nor the granting authority can be held responsible for them. Furkan Şahinuç is supported by the Konrad Zuse School of Excellence in Learning and Intelligent Systems (ELIZA) through the DAAD programme Konrad Zuse Schools of Excellence in Artificial Intelligence, sponsored by the Federal Ministry of Education and Research. We also thank Sheng Lu, Tim Baumgärtner (UKP Lab, TU Darmstadt), and Huiyin Xue (University of Sheffield) for their feedback during the writing of this paper.

## References

Osama Mohammed Afzal, Preslav Nakov, Tom Hope, and Iryna Gurevych. 2026. *Beyond “not novel enough”: Enriching scholarly critique with LLM-assisted feedback*. In *Proceedings of the 19th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2648–2671, Rabat, Morocco. Association for Computational Linguistics.

- Andrei Alexandru, Antonia Calvi, Henry Broomfield, Jackson Golden, Kyle Dai, Mathias Leys, Maurice Burger, Max Bartolo, Roman Engeler, Sashank Pisupati, Toby Drane, and Young Sun Park. 2025. *Atla selene mini: A general purpose evaluation model*. Preprint, arXiv:2501.17195.
- Zachary Ankner, Mansheej Paul, Brandon Cui, Jonathan D. Chang, and Prithviraj Ammanabrolu. 2024. *Critique-out-loud reward models*. Preprint, arXiv:2408.11791.
- Yuntao Bai, Saurav Kadavath, Sandipan Kundu, Amanda Askell, Jackson Kernion, Andy Jones, Anna Chen, Anna Goldie, Azalia Mirhoseini, Cameron McKinnon, Carol Chen, Catherine Olsson, Christopher Olah, Danny Hernandez, Dawn Drain, Deep Ganguli, Dustin Li, Eli Tran-Johnson, Ethan Perez, and 32 others. 2022. *Constitutional ai: Harmlessness from ai feedback*. Preprint, arXiv:2212.08073.
- Tuhin Chakrabarty, Philippe Laban, and Chien-Sheng Wu. 2025. *AI-slop to AI-polish? aligning language models through edit-based writing rewards and test-time computation*. In *Second Conference on Language Modeling*.
- Nuo Chen, Zhiyuan Hu, Qingyun Zou, Jiaying Wu, Qian Wang, Bryan Hooi, and Bingsheng He. 2025. *Judgelm: Large reasoning models as a judge*. Preprint, arXiv:2504.00050.
- Xiushi Chen, Gaotang Li, Ziqi Wang, Bowen Jin, Cheng Qian, Yu Wang, Hongru WANG, Yu Zhang, Denghui Zhang, Tong Zhang, Hanghang Tong, and Heng Ji. 2026. *RM-r1: Reward modeling as reasoning*. In *Fourteenth International Conference on Learning Representations*.
- Furkan Şahinuç, Subhabrata Dutta, and Iryna Gurevych. 2025. *Expert preference-based evaluation of automated related work generation*. Preprint, arXiv:2508.07955.
- Unsloth team Daniel Han, Michael Han. 2023. *Unsloth*.
- Flow AI. 2024. *Flow judge: An open small language model for llm system evaluations*. <https://flow-ai.com/blog/flow-judge>. Technical report, accessed: 2025-12-27.
- 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 *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*, 51:661–687.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, and 1 others. 2024. *The llama 3 herd of models*. Preprint, arXiv:2407.21783.
- Daya Guo, Dejian Yang, Haowei Zhang, and 1 others. 2025a. *Deepseek-r1 incentivizes reasoning in llms through reinforcement learning*. *Nature*, 645:633–638.
- Jiaxin Guo, Zewen Chi, Li Dong, Qingxiu Dong, Xun Wu, Shaohan Huang, and Furu Wei. 2025b. *Reward reasoning models*. In *Thirty-ninth Annual Conference on Neural Information Processing Systems*.
- Edward J Hu, yelong shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2022. *LoRA: Low-rank adaptation of large language models*. In *International Conference on Learning Representations*.
- IBM. 2025. *Granite-3.3-8b-instruct*. <https://huggingface.co/ibm-granite/granite-3.3-8b-instruct>. Accessed: 2025-12-11.
- Ai Jian, Jingqing Ruan, Xing Ma, Dailin Li, Qian-Lin Zhou, Ke Zeng, and Xunliang Cai. 2025. *Patarm: Bridging pairwise and pointwise signals via preference-aware task-adaptive reward modeling*. Preprint, arXiv:2510.24235.
- Leane Jourdan, Nicolas Hernandez, Florian Boudin, and Richard Dufour. 2025. *Identifying reliable evaluation metrics for scientific text revision*. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6731–6756, Vienna, Austria. Association for Computational Linguistics.
- Seungone Kim, Juyoung Suk, Shayne Longpre, Bill Yuchen Lin, Jamin Shin, Sean Welleck, Graham Neubig, Moontae Lee, Kyungjae Lee, and Minjoon Seo. 2024. *Prometheus 2: An open source language model specialized in evaluating other language models*. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 4334–4353, Miami, Florida, USA. Association for Computational Linguistics.
- Seungone Kim, Ian Wu, Jinu Lee, Xiang Yue, Seongyun Lee, Mingyeong Moon, Kiril Gashevski, Carolin Lawrence, Julia Hockenmaier, Graham Neubig, and Sean Welleck. 2025. *Scaling evaluation-time compute with reasoning models as process evaluators*. Preprint, arXiv:2503.19877.
- Xin Lai, Zhuotao Tian, Yukang Chen, Senqiao Yang, Xiangu Peng, and Jiaya Jia. 2024. *Step-dpo: Step-wise preference optimization for long-chain reasoning of llms*. Preprint, arXiv:2406.18629.
- Nathan Lambert, Valentina Pyatkin, Jacob Morrison, LJ Miranda, Bill Yuchen Lin, Khyathi Chandu, Nouha Dziri, Sachin Kumar, Tom Zick, Yejin Choi, Noah A. Smith, and Hannaneh Hajishirzi. 2025. *RewardBench: Evaluating reward models for language modeling*. In *Findings of the Association for Computational Linguistics: NAACL 2025*, pages 1755–1797, Albuquerque, New Mexico. Association for Computational Linguistics.

- Xiangci Li and Jessica Ouyang. 2024. [Related work and citation text generation: A survey](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 13846–13864, Miami, Florida, USA. Association for Computational Linguistics.
- Zhen Li, Xiaohan Xu, Tao Shen, Can Xu, Jia-Chen Gu, Yuxuan Lai, Chongyang Tao, and Shuai Ma. 2024. [Leveraging large language models for NLG evaluation: Advances and challenges](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 16028–16045, Miami, Florida, USA. Association for Computational Linguistics.
- Weixin Liang, Yuhui Zhang, Hancheng Cao, Binglu Wang, Daisy Yi Ding, Xinyu Yang, Kailas Vodrahalli, Siyu He, Daniel Scott Smith, Yian Yin, Daniel A. McFarland, and James Zou. 2024. [Can large language models provide useful feedback on research papers? a large-scale empirical analysis](#). *NEJM AI*, 1(8):AIoa2400196.
- Xiaobo Liang, Haoke Zhang, Juntao Li, Kehai Chen, Qiaoming Zhu, and Min Zhang. 2025. [Generative reward modeling via synthetic criteria preference learning](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 26755–26769, Vienna, Austria. Association for Computational Linguistics.
- Tianci Liu, Ran Xu, Tony Yu, Ilgee Hong, Carl Yang, Tuo Zhao, and Haoyu Wang. 2025a. [Openrubrics: Towards scalable synthetic rubric generation for reward modeling and llm alignment](#). *Preprint*, arXiv:2510.07743.
- Yang Liu, Dan Iter, Yichong Xu, Shuohang 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.
- 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](#). *Preprint*, arXiv:2504.02495.
- Qiyao Ma, Yunsheng Shi, Hongtao Tian, Chao Wang, Weiming Chang, and Ting Yao. 2025. [From faithfulness to correctness: Generative reward models that think critically](#). *Preprint*, arXiv:2509.25409.
- Saumya Malik, Valentina Pyatkin, Sander Land, Jacob Morrison, Noah A. Smith, Hannaneh Hajishirzi, and Nathan Lambert. 2026. [Reardeval: Advancing reward model evaluation](#). In *Fourteenth International Conference on Learning Representations*.
- OpenAI. 2025. Gpt-5.1: A smarter, more conversational chatgpt. <https://openai.com/index/gpt-5-1/>. Accessed: 2025-12-11.
- Sukannya Purkayastha, Zhuang Li, Anne Lauscher, Lizhen Qu, and Iryna Gurevych. 2025. [LazyReview: A dataset for uncovering lazy thinking in NLP peer reviews](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3280–3308, Vienna, Austria. Association for Computational Linguistics.
- Abdelrahman Sadallah, Tim Baumgärtner, Iryna Gurevych, and Ted Briscoe. 2025. [The good, the bad and the constructive: Automatically measuring peer review’s utility for authors](#). In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 28979–29009, Suzhou, China. Association for Computational Linguistics.
- Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, Y. K. Li, Y. Wu, and Daya Guo. 2024. [Deepseekmath: Pushing the limits of mathematical reasoning in open language models](#). *Preprint*, arXiv:2402.03300.
- Tu Shiwen, Zhao Liang, Chris Yuhao Liu, Liang Zeng, and Yang Liu. 2024. [Skywork critic model series](#). <https://huggingface.co/Skywork>.
- Joar Skalse, Nikolaus Howe, Dmitrii Krasheninnikov, and David Krueger. 2022. [Defining and characterizing reward gaming](#). In *Advances in Neural Information Processing Systems*, volume 35, pages 9460–9471. Curran Associates, Inc.
- Annalisa Szymanski, Noah Ziemis, Heather A. Eicher-Miller, Toby Jia-Jun Li, Meng Jiang, and Ronald A. Metoyer. 2025. [Limitations of the llm-as-a-judge approach for evaluating llm outputs in expert knowledge tasks](#). In *Proceedings of the 30th International Conference on Intelligent User Interfaces, IUI ’25*, page 952–966, New York, NY, USA. Association for Computing Machinery.
- Zecheng Tang, Baibei Ji, Quantong Qiu, Haitian Wang, Xiaobo Liang, Juntao Li, and Min Zhang. 2025. [Longrm: Revealing and unlocking the context boundary of reward modeling](#). *Preprint*, arXiv:2510.06915.
- Tianlu Wang, Iliia 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](#). *Preprint*, arXiv:2408.02666.
- Xumeng Wen, Zihan Liu, Shun Zheng, Shengyu Ye, Zhirong Wu, Yang Wang, Zhijian Xu, Xiao Liang, Junjie Li, Ziming Miao, Jiang Bian, and Mao Yang. 2026. [Reinforcement learning with verifiable rewards implicitly incentivizes correct reasoning in base LLMs](#). In *Fourteenth International Conference on Learning Representations*.
- Chenxi Whitehouse, Tianlu Wang, Ping Yu, Xian Li, Jason E Weston, Iliia Kulikov, and Swarnadeep Saha. 2026. [J1: Incentivizing thinking in LLM-as-a-judge via reinforcement learning](#). In *Fourteenth International Conference on Learning Representations*.

An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, Chujie Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng Hu, Hao Ge, Haoran Wei, Huan Lin, Jialong Tang, and 41 others. 2025a. [Qwen3 technical report](#). *Preprint*, arXiv:2505.09388.

An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, Huan Lin, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiayi Yang, Jingren Zhou, Junyang Lin, Kai Dang, and 23 others. 2025b. [Qwen2.5 technical report](#). *Preprint*, arXiv:2412.15115.

Rui Yang, Ruomeng Ding, Yong Lin, Huan Zhang, and Tong Zhang. 2024. [Regularizing hidden states enables learning generalizable reward model for llms](#). In *Advances in Neural Information Processing Systems*, volume 37, pages 62279–62309. Curran Associates, Inc.

Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhonghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, Hao Zhang, Joseph E Gonzalez, and Ion Stoica. 2023. [Judging llm-as-a-judge with mt-bench and chatbot arena](#). In *Advances in Neural Information Processing Systems*, volume 36, pages 46595–46623. Curran Associates, Inc.

## A Dataset

Since RevUtil dataset contains both human-annotated and synthetically generated labels, we merge these two splits. For the *Verifiability* aspect, we divide the evaluation into two sub-evaluations. We first determine whether the given review passage contains any claims under the *Verifiability extraction* evaluation. In the next step, we evaluate the degree of verifiability of data instances that accommodate claims. We provide the number of instances and the corresponding scoring rubrics for related work and RevUtil datasets in Tables 5 and 6, respectively. In the final dataset, we combine these two sources, resulting in 65,357 data instances, with 58,712 used for training and 6,645 reserved for testing.

Rel. Work	Train	Test	Scoring
Positioning Type	954	204	0–1
Positioning Cons.	2,822	605	0–1
Coherence	4,890	1,048	0–1
<b>Total</b>	<b>8,666</b>	<b>1,857</b>	<b>10,523</b>

Table 5: Statistics and scoring rubric of each aspect in related work dataset.

RevUtil	Train	Test	Scoring
Actionability	10,432	1,000	1–5
Grounding	10,431	1,000	1–5
Verifiability Ext.	10,430	1,000	0–1
Verifiability	8,323	788	1–5
Helpfulness	10,430	1,000	1–5
<b>Total</b>	<b>50,046</b>	<b>4,788</b>	<b>54,834</b>

Table 6: Statistics and scoring rubric of each aspect in RevUtil dataset.

## B Experimental Details

We use the Unsloth (Daniel Han, 2023) framework for LoRA optimization. We select LoRA rank as 64. In the roll-out phase, we utilize vLLM framework<sup>7</sup> while generating candidate responses. For each data instance, we sample 4 prompts via nucleus sampling with temperature of 1 and top-p of 0.95. Training is performed on a single NVIDIA A100 GPU (80 GB) and takes approximately three days.

## C Evaluation Prompts

We provide the list of prompts used while evaluating scientific writings in Table 7.

Evaluation Type	Prompts
System prompt	Figure 5
RW - Coherence	Figures 6, 7
RW - Pos. Type	Figures 8, 9, 10, 11
RW - Pos. Cons.	Figures 12, 13, 14, 15
RU - Verifiability Ext.	Figure 16
RU - Verifiability	Figures 17, 18
RU - Actionability	Figures 19, 20
RU - Grounding Spec.	Figures 21, 22
RU - Helpfulness	Figures 23, 24
Nov. Alignment	Figures 25, 26
Rev. - Relatedness	Figures 27, 28
Rev. - Correctness	Figures 29, 30

Table 7: Evaluation prompts table

## D Full Results

We provide full numeric results in Tables 8, 9, 10 for accuracy and 11, 12, 13 for F1 Score.

## E Reasoning Trace Analysis

We provide full reasoning trace examples for review utility, novelty alignment evaluation and re-

<sup>7</sup><https://docs.vllm.ai/en/stable/>

<b>Rel. Work</b>	<b>Coherence</b>	<b>Positioning Cons.</b>	<b>Positioning Type</b>	<b>Overall Task</b>
DS-Qwen	0.61 ± 0.010	0.85 ± 0.010	0.66 ± 0.017	0.70 ± 0.007
Granite3.3	0.66 ± 0.007	0.88 ± 0.006	0.63 ± 0.016	0.73 ± 0.005
Llama3.1	0.62 ± 0.013	0.89 ± 0.010	0.60 ± 0.029	0.70 ± 0.012
Qwen2.5	0.67 ± 0.010	0.94 ± 0.006	0.68 ± 0.018	0.76 ± 0.008
Qwen3	0.65 ± 0.003	0.95 ± 0.009	0.96 ± 0.009	0.78 ± 0.003
Prometheus	0.64 ± 0.012	0.87 ± 0.011	0.61 ± 0.017	0.71 ± 0.007
Selene	0.63 ± 0.010	0.88 ± 0.004	0.59 ± 0.016	0.71 ± 0.005
Skywork	0.69 ± 0.005	0.94 ± 0.006	0.78 ± 0.030	0.78 ± 0.005
DS-GRM	0.57 ± 0.012	0.85 ± 0.012	0.71 ± 0.011	0.68 ± 0.010
o3-mini	0.63 ± 0.003	0.99 ± 0.002	0.998 ± 0.002	0.79 ± 0.002
GPT-5.2	0.73 ± 0.002	0.94 ± 0.004	1.00 ± 0.000	0.83 ± 0.002
Sci-RM	0.73 ± 0.007	0.995 ± 0.001	0.86 ± 0.011	0.83 ± 0.005
Sci-RM Ref.	0.73 ± 0.004	0.996 ± 0.001	0.87 ± 0.012	0.83 ± 0.003

Table 8: Full accuracy results for related work evaluation. Values are reported as mean ± standard deviation.

<b>RevUtil</b>	<b>Verifiability Ext.</b>	<b>Verifiability</b>	<b>Actionability</b>	<b>Grounding Spec.</b>	<b>Helpfulness</b>	<b>Overall Task</b>
DS-Qwen	0.81 ± 0.005	0.39 ± 0.017	0.30 ± 0.006	0.34 ± 0.004	0.36 ± 0.014	0.44 ± 0.003
Granite3.3	0.82 ± 0.006	0.39 ± 0.011	0.32 ± 0.004	0.40 ± 0.011	0.36 ± 0.009	0.46 ± 0.005
Llama3.1	0.87 ± 0.004	0.27 ± 0.006	0.30 ± 0.009	0.43 ± 0.015	0.42 ± 0.007	0.47 ± 0.004
Qwen2.5	0.85 ± 0.007	0.38 ± 0.016	0.32 ± 0.008	0.40 ± 0.006	0.45 ± 0.014	0.49 ± 0.005
Qwen3	0.88 ± 0.005	0.50 ± 0.012	0.42 ± 0.010	0.60 ± 0.005	0.43 ± 0.010	0.57 ± 0.005
Prometheus	0.80 ± 0.015	0.40 ± 0.006	0.36 ± 0.017	0.29 ± 0.013	0.40 ± 0.008	0.45 ± 0.010
Selene	0.87 ± 0.003	0.44 ± 0.005	0.32 ± 0.015	0.40 ± 0.005	0.46 ± 0.010	0.50 ± 0.005
Skywork	0.90 ± 0.006	0.42 ± 0.019	0.43 ± 0.010	0.48 ± 0.011	0.42 ± 0.010	0.53 ± 0.005
DS-GRM	0.70 ± 0.009	0.39 ± 0.011	0.35 ± 0.015	0.39 ± 0.010	0.41 ± 0.008	0.45 ± 0.005
o3-mini	0.86 ± 0.002	0.48 ± 0.007	0.49 ± 0.010	0.63 ± 0.007	0.51 ± 0.005	0.60 ± 0.001
GPT-5.2	0.89 ± 0.002	0.46 ± 0.004	0.40 ± 0.008	0.66 ± 0.004	0.49 ± 0.009	0.59 ± 0.004
Sci-RM	0.92 ± 0.002	0.65 ± 0.007	0.62 ± 0.009	0.70 ± 0.006	0.63 ± 0.012	0.71 ± 0.006
Sci-RM Ref.	0.92 ± 0.005	0.65 ± 0.007	0.58 ± 0.007	0.70 ± 0.008	0.62 ± 0.003	0.69 ± 0.003

Table 9: Full accuracy results for review utility evaluation. Values are reported as mean ± standard deviation.

sion evaluation across different models in Tables 14, 15, and 16, respectively.

Model	Novelty	Revision		
	Alignment	Relatedness	Correctness	Overall Task
DS-Qwen	0.41 ± 0.07	0.77 ± 0.001	0.67 ± 0.005	0.72 ± 0.002
Granite3.3	0.48 ± 0.04	0.80 ± 0.003	0.68 ± 0.003	0.74 ± 0.003
Llama3.1	0.56 ± 0.04	0.56 ± 0.009	0.61 ± 0.006	0.58 ± 0.005
Qwen2.5	0.71 ± 0.05	0.80 ± 0.005	0.66 ± 0.010	0.73 ± 0.006
Qwen3	0.58 ± 0.05	0.86 ± 0.004	0.76 ± 0.003	0.81 ± 0.003
Prometheus	0.33 ± 0.06	0.76 ± 0.006	0.70 ± 0.003	0.73 ± 0.002
Selene	0.57 ± 0.04	0.79 ± 0.004	0.71 ± 0.004	0.75 ± 0.003
Skywork	0.42 ± 0.05	0.83 ± 0.003	0.77 ± 0.004	0.80 ± 0.002
DS-GRM	0.35 ± 0.05	0.73 ± 0.001	0.59 ± 0.008	0.66 ± 0.004
o3-mini	0.78 ± 0.01	0.88 ± 0.002	0.80 ± 0.002	0.84 ± 0.0008
GPT-5.2	0.74 ± 0.06	0.82 ± 0.003	0.78 ± 0.002	0.80 ± 0.001
Sci-RM	0.61 ± 0.04	0.82 ± 0.001	0.73 ± 0.005	0.78 ± 0.002
Sci-RM Ref.	0.74 ± 0.04	0.82 ± 0.002	0.73 ± 0.003	0.78 ± 0.002

Table 10: Full accuracy results for unseen novelty and revision evaluation tasks. Values are reported as mean ± standard deviation.

Rel. Work	Coherence	Positioning Cons.	Positioning Type	Overall Task
DS-Qwen	0.61 ± 0.012	0.85 ± 0.011	0.66 ± 0.019	0.71 ± 0.007
Granite3.3	0.66 ± 0.009	0.88 ± 0.007	0.64 ± 0.019	0.73 ± 0.006
Llama3.1	0.63 ± 0.016	0.89 ± 0.011	0.61 ± 0.031	0.71 ± 0.016
Qwen2.5	0.67 ± 0.012	0.95 ± 0.005	0.70 ± 0.022	0.77 ± 0.009
Qwen3	0.65 ± 0.001	0.95 ± 0.011	0.96 ± 0.008	0.86 ± 0.003
Prometheus	0.64 ± 0.014	0.88 ± 0.013	0.63 ± 0.017	0.72 ± 0.008
Selene	0.63 ± 0.011	0.88 ± 0.004	0.58 ± 0.017	0.70 ± 0.003
Skywork	0.69 ± 0.006	0.94 ± 0.007	0.77 ± 0.036	0.80 ± 0.013
DS-GRM	0.58 ± 0.014	0.86 ± 0.011	0.68 ± 0.019	0.71 ± 0.010
o3-mini	0.59 ± 0.003	0.99 ± 0.002	0.998 ± 0.003	0.86 ± 0.001
GPT-5.2	0.72 ± 0.003	0.94 ± 0.005	1.00 ± 0.000	0.89 ± 0.002
Sci-RM	0.73 ± 0.008	0.99 ± 0.001	0.86 ± 0.013	0.86 ± 0.005
Sci-RM Ref.	0.73 ± 0.004	0.996 ± 0.001	0.87 ± 0.013	0.87 ± 0.005

Table 11: Full F1 score results for related work evaluation. Values are reported as mean ± standard deviation.

RevUtil	Verifiability Ext.	Verifiability	Actionability	Grounding Spec.	Helpfulness	Overall Task
DS-Qwen	0.82 ± 0.004	0.42 ± 0.018	0.29 ± 0.006	0.35 ± 0.004	0.38 ± 0.015	0.45 ± 0.004
Granite3.3	0.80 ± 0.007	0.41 ± 0.014	0.34 ± 0.003	0.42 ± 0.014	0.37 ± 0.010	0.47 ± 0.007
Llama3.1	0.87 ± 0.005	0.31 ± 0.010	0.31 ± 0.015	0.41 ± 0.021	0.43 ± 0.009	0.47 ± 0.005
Qwen2.5	0.85 ± 0.008	0.41 ± 0.017	0.32 ± 0.011	0.39 ± 0.009	0.42 ± 0.015	0.48 ± 0.006
Qwen3	0.88 ± 0.005	0.52 ± 0.012	0.44 ± 0.012	0.60 ± 0.006	0.44 ± 0.011	0.57 ± 0.005
Prometheus	0.81 ± 0.016	0.42 ± 0.008	0.36 ± 0.018	0.34 ± 0.012	0.41 ± 0.011	0.47 ± 0.009
Selene	0.85 ± 0.005	0.45 ± 0.003	0.26 ± 0.014	0.42 ± 0.006	0.43 ± 0.010	0.48 ± 0.004
Skywork	0.89 ± 0.008	0.44 ± 0.024	0.43 ± 0.011	0.48 ± 0.015	0.44 ± 0.013	0.53 ± 0.007
DS-GRM	0.73 ± 0.008	0.40 ± 0.012	0.33 ± 0.019	0.41 ± 0.010	0.41 ± 0.007	0.46 ± 0.006
o3-mini	0.83 ± 0.005	0.51 ± 0.010	0.48 ± 0.016	0.61 ± 0.008	0.52 ± 0.008	0.59 ± 0.002
GPT-5.2	0.88 ± 0.004	0.47 ± 0.009	0.37 ± 0.010	0.65 ± 0.005	0.49 ± 0.011	0.57 ± 0.006
Sci-RM	0.92 ± 0.003	0.59 ± 0.010	0.59 ± 0.010	0.67 ± 0.007	0.62 ± 0.014	0.68 ± 0.007
Sci-RM Ref.	0.92 ± 0.006	0.59 ± 0.010	0.53 ± 0.008	0.66 ± 0.008	0.60 ± 0.004	0.66 ± 0.004

Table 12: Full F1 score results for review utility evaluation. Values are reported as mean ± standard deviation.

Model	Novelty	Revision		
	Alignment	Relatedness	Correctness	Overall Task
DS-Qwen	0.50 ± 0.092	0.74 ± 0.002	0.61 ± 0.008	0.67 ± 0.004
Granite3.3	0.51 ± 0.041	0.75 ± 0.006	0.61 ± 0.005	0.68 ± 0.005
Llama3.1	0.61 ± 0.043	0.61 ± 0.009	0.61 ± 0.007	0.61 ± 0.005
Qwen2.5	0.69 ± 0.063	0.79 ± 0.006	0.65 ± 0.007	0.72 ± 0.005
Qwen3	0.61 ± 0.054	0.86 ± 0.006	0.74 ± 0.004	0.80 ± 0.004
Prometheus	0.39 ± 0.071	0.78 ± 0.004	0.69 ± 0.004	0.74 ± 0.002
Selene	0.62 ± 0.046	0.79 ± 0.004	0.70 ± 0.005	0.75 ± 0.003
Skywork	0.43 ± 0.060	0.83 ± 0.003	0.77 ± 0.005	0.80 ± 0.002
DS-GRM	0.30 ± 0.091	0.76 ± 0.001	0.61 ± 0.009	0.69 ± 0.004
o3-mini	0.78 ± 0.010	0.87 ± 0.002	0.79 ± 0.002	0.83 ± 0.001
GPT-5.2	0.74 ± 0.010	0.83 ± 0.004	0.78 ± 0.001	0.80 ± 0.001
Sci-RM	0.63 ± 0.046	0.81 ± 0.002	0.69 ± 0.005	0.75 ± 0.003
Sci-RM Ref.	0.70 ± 0.040	0.81 ± 0.002	0.69 ± 0.004	0.75 ± 0.003

Table 13: Full F1 score results for unseen novelty and revision evaluation tasks. Values are reported as mean ± standard deviation.

Task: Review Utility	Aspect: Actionability
<b>Query Review:</b> Evaluation datasets are all short outputs: classification or question answering. This is a slightly limited setting for the claims.	
<b>Gold label:</b> 1	
<b>Label 1 Explanation:</b> The comment lacks meaningful information to help authors improve the paper. Authors do not know what they should do after reading the comment.	
<b>Label 3 Explanation:</b> The comment explicitly states an action but is vague on how to execute it.	
SciRM	Qwen2.5-7B-Instruct
<p>&lt;reasoning&gt;The provided comment, "Evaluation datasets are all short outputs: classification or question answering. This is a slightly limited setting for the claims," lacks specific and actionable guidance. It points out a limitation but does not offer a concrete solution or modification. The authors are left to determine how to expand the evaluation datasets, but without more detail, this action is too vague to be considered as a clear and actionable instruction. As a result, this comment is not sufficient for an actionable evaluation score. Therefore, the evaluation score should be 1, indicating that the comment lacks meaningful information to help the authors improve the paper. The reviewer is likely highlighting a limitation and expecting authors to address it, but without providing a clear direction, the comment is not actionable. The provided example would be improved by specifying particular actions or concrete steps that could be taken to improve the dataset or strengthen the claims in the paper.&lt;/reasoning&gt;</p> <p>&lt;score&gt;1&lt;/score&gt;</p>	<p>&lt;reasoning&gt;The comment suggests that the evaluation datasets used in the paper are predominantly short outputs for classification or question answering, implying a limited setting for the claims. It identifies a limitation but doesn't provide specific directions or a clear path for addressing this issue, making it somewhat vague. The authors might infer that they need to use a wider range of datasets, but the comment does not specify how they should go about doing so. Hence, while it is concrete in identifying the limitation, it lacks explicit instructions on how to solve it, resulting in a moderate level of actionability. Therefore, the appropriate score is 3.&lt;/reasoning&gt;</p> <p>&lt;score&gt;3&lt;/score&gt;</p>

Table 14: Reasoning trace comparison for SciRM and base model Qwen2.5-7B-Instruct

---

**Task:** Novelty Alignment**Aspect:** Coherence

---

**Query Assessments:** Assessment 1: (...) While the approach leverages transfer learning between 1D and 3D representations and incorporates enhancements in model scaling and architecture, the novelty from a machine learning perspective is limited. The core methodology primarily involves integrating existing components—specifically, LLaMA-based language models and established diffusion models—without introducing fundamentally new techniques. (...) Overall, the contribution is largely an application and combination of known methods rather than a significant advancement in machine learning methodology (...)

Assessment 2: NExT-Mol introduces a novel approach by explicitly transferring knowledge from a large-scale 1D molecular language model (SELFIES-based MoLlama) to a 3D diffusion model for conformer generation, using a cross-modal projector—a combination not previously demonstrated in the literature. While prior works have explored 1D language modeling and 3D diffusion separately or in joint 2D/3D settings, NExT-Mol is the first to leverage large-scale 1D pretraining for 3D molecular generation, resulting in incremental but real improvements on standard 3D benchmarks (FCD, COV-R). The claimed architectural innovations may represent refinements rather than conceptual leaps, and some terminology (e.g., "cross-modal projector") repackages established techniques. The authors' claims of 100% validity and transformative novelty are somewhat overstated, as 100% validity is a property of SELFIES and prior models have also scaled up or improved validity. Overall, the main contribution is a substantive integration of established paradigms (large-scale 1D LMs and 3D diffusion) via transfer learning, with advances that are meaningful but largely incremental in the context of a rapidly evolving field.

**Gold label:** 0

**Label 0 Explanation:** Two assessments do not come to the same conclusion in terms of novelty verdict.

**Label 1 Explanation:** Two assessments come to the same conclusion in terms of novelty verdict.

---

**SciRM-Ref**

<reasoning>Both assessments emphasize the integration of existing techniques but present different conclusions. Assessment 1 concludes the contribution is not novel, focusing on the lack of fundamental new methodologies and minimal innovation. In contrast, Assessment 2 calls the contribution novel, highlighting the combination of large-scale 1D pretraining for 3D molecular generation as an incremental but real improvement, though still with some limitations in novelty as argued. For the final evaluation, the key is whether both assessments come to the same conclusion regarding the novelty, and in this case, they do not. Assessment 1 considers the contribution non-novel, while Assessment 2 leans towards it being novel but with limited novelty. Therefore, the score should be 0, indicating that the assessments do not align in their final conclusions on the novelty of the paper. Both assessments provide detailed reasoning for their conclusions, with one suggesting the work lacks fundamental novelty and the other claiming it is an incremental but still novel contribution. The scoring rubric requires a binary judgment on whether the assessments come to the same conclusion about the novelty. Here, given their differing judgments, the score should be 0.</reasoning> <score>0</score>

**SciRM**

<reasoning>In the first assessment, the paper is considered not novel as the integration of existing models is deemed minimal and not providing significant advancement. The second assessment also indicates novelty, but recognizes a real, incremental improvement from prior work. Both assessments explicitly argue for a certain level of novelty, but they converge in agreeing that the contributions are not fundamental or groundbreaking, rather they represent a meaningful advancement but are incremental. Hence, the assessments come to the same conclusion: the work is not fundamentally novel, but it does offer incremental but significant contributions. Therefore, a score of 1 is appropriate.</reasoning> <score>1</score>

---

Table 15: Reasoning trace comparison for SciRM-Ref and SciRM

Task: Revision Evaluation	Aspect: Correctness
<p><b>Query Revision:</b> [ORIGINAL TEXT]: Before we define the diffusion distance, we briefly introduce the intuition behind it: two nodes are considered similar when they diffuse in a similar way through the graph, and therefore when they influence the other nodes in a similar manner (Fouss et al., 2012). In other words, two nodes are close if they are in the same cluster which has a consistent local structure. More precisely, the diffusion distance at time K between nodes i and j is defined as follows:</p> <p>[INSTRUCTION]: Remove the ideas which are not particularly essential for the overall paragraph.</p> <p>[ANSWER]: Before we define the diffusion distance, we briefly introduce the intuition behind it: two nodes are considered similar when they diffuse in a similar way through the graph, and therefore when they influence the other nodes in a similar manner (Fouss et al., 2012). More precisely, the diffusion distance at time K between nodes i and j is defined as follows:</p> <p><b>Gold label: 1</b></p> <p><b>Label 0 Explanation:</b> The revision is not a better version that can replace the original paragraph in the scientific article.</p> <p><b>Label 1 Explanation:</b> The revision is a better version that should replace the original paragraph in the scientific article.</p>	
<p><b>SciRM-Ref</b></p> <p>&lt;reasoning&gt;The instruction asks to remove ideas that are not essential for the overall paragraph. The answer successfully removes the phrase "In other words, two nodes are close if they are in the same cluster which has a consistent local structure," which is not as essential to the definition of diffusion distance as the other parts of the text. The resulting text maintains the essential information for the definition of diffusion distance and is a better version of the paragraph. Therefore, the evaluation score should be 1, indicating that the revision is a better version that could replace the original paragraph. In this case, the revised text focuses on the key points of the paragraph without any non-essential ideas, making it a stronger version that meets the criteria. This removal leads to a clearer, more concise definition without losing any critical information, ensuring that the revised version is a good replacement for the original.&lt;/reasoning&gt; &lt;score&gt;1&lt;/score&gt;</p>	<p><b>DS-GRM</b></p> <p>&lt;reasoning&gt;The instruction asks to remove the ideas which are not particularly essential for the overall paragraph. The provided answer keeps the main idea of the paragraph, which is introducing the diffusion distance and its relation to the similarity of nodes. However, it does not remove any non-essential details. The revised text is still informative and relevant, but it could be more concise by removing the phrase "in other words" and the subsequent explanation about the local structure of clusters, which are not particularly essential for the overall paragraph. Therefore, the evaluation score should be 0.&lt;/reasoning&gt; &lt;score&gt;0&lt;/score&gt;</p>

Table 16: Reasoning trace comparison for SciRM-Ref and DS-GRM baseline

**System Prompt**

You are an evaluator of expert-domain scientific writing. You will get a query-answer pair along with criteria explaining the specific evaluation aspect and the scoring rubric. You should evaluate whether the answer satisfy the query based on the given criteria. In addition, examples demonstrating how the evaluation should be performed will be provided. First output your reasoning enclosed between <reasoning> and </reasoning>. Then output your score enclosed between <score> and </score>. Inside <score> provide only the numeric score and nothing else.

Figure 5: System prompt that is commonly used across each evaluation task and aspect.

### Prompt for Related Work Evaluation - Coherence (Part 1)

**[QUERY]:** Given a context from a scientific paper, your task is to write a coherent citation sentence that cites the given paper with a specific citation number. Citation number will be also provided along with the context.

**[CRITERIA]:** The paper context supports (entails) the citation sentence. In cases where more than one paper is referenced in the sentence, as long as context in which given citation number fits the paper content, it should be count as entailment as well. In multiple citation cases, the paper does not have to entail whole sentence. Scoring rubric is as follows:

0: The given context does not align with the context in which the citation number appears in the sentence, or with the overall meaning of the sentence itself.

1: The given context is compatible with the context in which the citation number appears in the sentence.

**[EXAMPLES]:**

<START OF EXAMPLE 1>

**CONTEXT:** {*Abstract and Introduction of example cited paper*}

**CITATION NUMBER:** 10

**ANSWER:** Additionally, simple predictors that regress on features extracted from architectures, such as textual encoding schemes or direct structural descriptions, have been found effective for performance approximation, enabling significant sample efficiency improvements [9][10].

**EVALUATION:** <reasoning>The context discusses the effectiveness of a Neural Predictor method that uses regression modeling to predict validation accuracy based on architecture features, achieving significant sample efficiency improvements over traditional approaches. The citation sentence aligns with the context in pointing out the effectiveness of simple predictors for performance approximation, which correlates with the regression steps described in the Neural Predictor method. Therefore, the citation sentence is entailed by the context and follows from it. Therefore, the evaluation score should be 1.</reasoning><score>1</score>

<END OF EXAMPLE 1>

Figure 6: Prompt for Related Work Evaluation - Coherence (Part 1)

### Prompt for Related Work Evaluation - Coherence (Part 2)

<START OF EXAMPLE 2>

**CONTEXT:** {*Abstract and Introduction of example cited paper*}

**CITATION NUMBER:** 11

**ANSWER:** Natural Language Inference has emerged as an essential tool for modeling semantic relationships, with applications extending from textual entailment assessments [11,12] to dialogue consistency verification [17].

**EVALUATION:** <reasoning>The citation sentence discusses Natural Language Inference (NLI) and its applications, which is not directly related or entailed by the presented context on dialogue generation models and their improvements. The context does not specifically mention or address NLI or its use in dialogue consistency verification. As such, the hypothesis cannot be inferred or entailed from the given context. Therefore, the evaluation score should be 0.</reasoning><score>0</score>

<END OF EXAMPLE 2>

**[CONTEXT]:** {*Abstract and Introduction of corresponding cited paper*}

**[ANSWER]:** {*Data instance to be evaluated*}

Figure 7: Prompt for Related Work Evaluation - Coherence (Part 2)

### Prompt for Related Work Evaluation - Positioning Type (Each Paragraph) (Part 1)

**[QUERY]:** Your task is to write related work section for a scientific paper that states the paper's contribution or its position among the literature in each paragraph.

**[CRITERIA]:** In a multiple paragraph related work section, every single paragraph should state the contribution of the paper and/or its position among the literature. Ensure that contribution and/or positioning statements should be aligned with the specific focus of each paragraph. Scoring rubric is as follows:

0: The answer fails to state contribution of the paper and/or its position among the literature in each related work paragraph.

1: The answer states contribution of the paper and/or its position among the literature in each related work paragraph.

**[EXAMPLES]:**

<START OF EXAMPLE 1>

**ANSWER:** Previous studies on TDMR (Task-Dataset-Metric-Result) extraction primarily rely on data sources such as NLP-progress or PapersWithCode. While these sources are popular, they lack rigorous quality control, including standardized TDM entity representation and complete publication coverage. Some efforts, such as [1] and [2], go further by extracting TDM triples and result values, and normalizing them for leaderboard construction. However, these approaches are constrained by their reliance on a closed domain, requiring extracted triples to match a pre-defined TDM taxonomy. Other studies, such as [3] and [4], only extract TDM triples without result values, thereby forgoing leaderboard construction. Meanwhile, [6] focuses on extracting result values, but only in the context of pre-defined TDM triples, making it similar in limitation to [1]. Consequently, these methods are not adaptable to the dynamic nature of emerging benchmarks, where new tasks, datasets, and metrics are frequently introduced.

In a parallel research thread, scientific leaderboards have been treated as a form of scientific knowledge graphs composed of Task, Dataset, and Metric entities. Information extraction in this context has been widely studied, as seen in works such as [5, 7, 8, 9, 10]. These studies typically focus on linking scientific content to structured entities but often assume that the schema or triples are known in advance.

In contrast to prior work, our approach addresses several critical limitations. We construct our dataset manually from publications to ensure comprehensive TDMR annotations, rather than relying on incomplete or unstandardized third-party sources. We normalize TDMR tuples for accurate leaderboard construction and propose novel experimental settings that reflect realistic, dynamic scenarios, including cold-start conditions without any pre-defined TDM taxonomy. By simultaneously constructing a scientific knowledge graph and linking publications to it in an open-domain setting, our method is inherently adaptable to the evolving research landscape.

**EVALUATION:** <reasoning>The related work draft does not mention the main paper's contributions and its position in each individual paragraph; instead, it waits until the final paragraph to outline the specific contributions of the study. The contributions are clearly summarized and detailed only in the final paragraph. Therefore, the evaluation score should be 0.</reasoning><score>0</score>

<END OF EXAMPLE 1>

Figure 8: Prompt for Related Work Evaluation - Positioning Type (Each Paragraph) (Part 1)

### Prompt for Related Work Evaluation - Positioning Type (Each Paragraph) (Part 2)

<START OF EXAMPLE 2>

**ANSWER:** In terms of data source, previous studies use either NLP-progress or paperswithcode. These sources, however, lack rigorous quality assurance, such as standardizing TDM entities across different leaderboards and ensuring complete coverage of relevant publications. Similar to our work, [1] and [2] extract TDM triples along with the results values and apply normalization for leaderboard construction. However, both studies assume a closed domain and match extracted TDM triples to a pre-defined TDM triple taxonomy. On the other hand, some studies only partially extract TDMR tuples and do not apply normalization. For example, [3] and [4] extract TDM triples without results. Therefore, these works do not deal with leaderboard construction. In addition, [6] extract the results values depending on the pre-defined TDM triples. Both [3] and [6] leverage pre-defined TDM triples in an extraction process similar to [1]. Since these approaches require a pre-defined taxonomy of TDM triples, they are incompatible with a realistic task definition. In short, none of the previous work is adaptable to the constantly emerging benchmarks driven by new research and innovation. In this work, we address the aforementioned problems. Unlike previous work, we (1) manually construct our dataset directly from publications to ensure complete TDMR annotations, (2) apply normalization for leaderboard construction, and (3) propose different experimental settings to simulate real-world scenarios.

Part of the scientific leaderboards can be viewed as a special type of scientific knowledge graph that includes three types of entities (Task, Dataset, Metric) and the relations between them, which have been the primary focus of the previous studies on information extraction from scientific literature [5, 7, 8, 9, 10]. Our work in the cold start scenario, in which we do not assume any pre-defined TDM triple is given, constructs such a scientific knowledge graph and links the papers to the nodes in the graph simultaneously.

**EVALUATION:** <reasoning>Contributions and position of the paper in the literature are mentioned within each paragraph rather than being summarized at the end. The first paragraph discusses previous work and their limitations, then explicitly states what the current work does differently, listing three specific contributions. The second paragraph again discusses previous studies but also explains how the current work differs by constructing a knowledge graph without assuming predefined triples and linking papers to nodes. Therefore, the evaluation score should be 1.</reasoning><score>1</score>

<END OF EXAMPLE 2>

Figure 9: Prompt for Related Work Evaluation - Positioning Type (Each Paragraph) (Part 2)

### Prompt for Related Work Evaluation - Positioning Type (Each Paragraph) (Part 3)

<START OF EXAMPLE 3>

**ANSWER:** Several efforts have been made to automate the extraction and organization of performance results from scientific publications. Early work by [1] introduced a framework, TDMS-IE, for identifying task, dataset, metric, and score tuples from NLP papers, establishing a foundation for automatic leaderboard construction. Similarly, AxCell [2] proposes a robust pipeline that utilizes structural analysis and novel table segmentation techniques to extract results from machine learning papers, demonstrating significant improvements over prior methods. ORKG-Leaderboards [3] offers a systematic approach that integrates leaderboard extraction into a knowledge graph framework, enabling machine-actionable publishing and dynamic visualization of state-of-the-art performance. TELIN [4] focuses on extracting leaderboard-relevant entities from PDFs using a semi-automated approach that reduces human annotation needs through targeted entity refinement.

In addition to these extraction systems, several datasets and benchmarks have been introduced to facilitate the development and evaluation of leaderboard construction tools. LEGOBench [5] provides a large-scale benchmark derived from arXiv and PapersWithCode, and evaluates both language model-based and graph-based approaches. SciERC and its associated framework SciIE [6] support multi-task extraction of entities and relations, enabling construction of scientific knowledge graphs. SciREX [7] extends information extraction to the document level, capturing relationships that span across sections, which is critical for leaderboard generation from full papers.

Other contributions include TDMSci [8], a specialized corpus annotated with task, dataset, and metric entities, which supports the development of more accurate extraction models. SciNLP-KG [9] introduces methods for extracting entity relations from NLP literature to build a knowledge graph, highlighting its potential use in leaderboard automation. Additionally, a diachronic analysis of NLP research trends by [10] shows the evolving influence of tasks, methods, and datasets, providing a broader context for understanding the dynamic nature of scientific benchmarks.

**EVALUATION:** <reasoning>The related work section draft provided does not explicitly mention or discuss the main paper's contributions or its position. Each paragraph focuses on summarizing existing research efforts and methodologies without indicating how the current paper builds upon or differs from these works. Additionally, the final paragraph does not serve as a summary of the main paper's contributions; instead, it continues to discuss other related works without tying them back to the current study's advancements. Therefore, the evaluation score should be 0.</reasoning><score>0</score>

<END OF EXAMPLE 3>

[ANSWER]: {*Data instance to be evaluated*}

Figure 10: Prompt for Related Work Evaluation - Positioning Type (Each Paragraph) (Part 3)

### Prompt for Related Work Evaluation - Positioning Type (Final Paragraph)

**[QUERY]:** Your task is to write related work section that contributions and/or positioning of the paper is provided in the final paragraph as a summary.

**[CRITERIA]:** Each paragraph (except final paragraph) should focus on summarizing and comparing prior research etc., without mentioning the current paper's contributions or its position in the literature. Then, the final paragraph at the end of the related work section should the current paper's contributions and its positioning. The final paragraph should address aspects mentioned in previous paragraphs. Scoring rubric is as follows:

0: The answer provide contributions in each paragraph or final paragraph fails to address the points in previous paragraphs or do not mention contributions-positioning at all.

1: Only the final paragraph emphasizes the contributions-positioning by addressing the points in previous paragraphs.

**[EXAMPLES]:**

<START OF EXAMPLE 1>

**ANSWER:** {*Example answer in positioning type (Each Paragraph) (Part 1)*}

**EVALUATION:** <reasoning>The related work draft does not mention the main paper's contributions and its position in each individual paragraph; instead, it waits until the final paragraph to outline the specific contributions of the study. The contributions are clearly summarized and detailed only in the final paragraph. Therefore, the evaluation score should be 1.</reasoning><score>1</score>

<END OF EXAMPLE 1>

<START OF EXAMPLE 2>

**ANSWER:** {*Example answer in positioning type (Each Paragraph) (Part 2)*}

**EVALUATION:** <reasoning>Contributions and position of the paper in the literature are mentioned within each paragraph rather than being summarized at the end. The first paragraph discusses previous work and their limitations, then explicitly states what the current work does differently, listing three specific contributions. The second paragraph again discusses previous studies but also explains how the current work differs by constructing a knowledge graph without assuming predefined triples and linking papers to nodes. Therefore, the evaluation score should be 0.</reasoning><score>0</score>

<END OF EXAMPLE 2>

<START OF EXAMPLE 3>

**ANSWER:** {*Example answer in positioning type (Each Paragraph) (Part 3)*}

**EVALUATION:** <reasoning>The related work section draft provided does not explicitly mention or discuss the main paper's contributions or its position. Each paragraph focuses on summarizing existing research efforts and methodologies without indicating how the current paper builds upon or differs from these works. Additionally, the final paragraph does not serve as a summary of the main paper's contributions; instead, it continues to discuss other related works without tying them back to the current study's advancements. Therefore, the evaluation score should be 0.</reasoning><score>0</score>

<END OF EXAMPLE 3>

**[ANSWER]:** {*Data instance to be evaluated*}

Figure 11: Prompt for Related Work Evaluation - Positioning Type (Final Paragraph)

### Prompt for Related Work Evaluation - Positioning Consistency (Each Paragraph) (Part 1)

**[QUERY]:** Your task is to write a related work section paragraph for a scientific paper that states the paper's contribution or its position among the literature.

**[CRITERIA]:** The generated paragraph should explicitly or implicitly mention the main paper's contribution or position among existing literature. Ensure that contribution and/or positioning statements should be aligned with the specific focus of the paragraph. Scoring rubric is as follows:

0: The answer paragraph fails to explicitly or implicitly mention the main paper's contribution or position among existing literature.

1: The answer paragraph explicitly or implicitly mention the main paper's contribution or position among existing literature.

**[EXAMPLES]:**

<START OF EXAMPLE 1>

**ANSWER:** In terms of data source, previous studies use either NLP-progress or paperswithcode. These sources, however, lack rigorous quality assurance, such as standardizing TDM entities across different leaderboards and ensuring complete coverage of relevant publications. Similar to our work, [1] and [2] extract TDM triples along with the results values and apply normalization for leaderboard construction. However, both studies assume a closed domain and match extracted TDM triples to a pre-defined TDM triple taxonomy. On the other hand, some studies only partially extract TDMR tuples and do not apply normalization. For example, [3] and [4] extract TDM triples without results. Therefore, these works do not deal with leaderboard construction. In addition, [6] extract the results values depending on the pre-defined TDM triples. Both [3] and [6] leverage pre-defined TDM triples in an extraction process similar to [1]. Since these approaches require a pre-defined taxonomy of TDM triples, they are incompatible with a realistic task definition. In short, none of the previous work is adaptable to the constantly emerging benchmarks driven by new research and innovation. In this work, we address the aforementioned problems. Unlike previous work, we (1) manually construct our dataset directly from publications to ensure complete TDMR annotations, (2) apply normalization for leaderboard construction, and (3) propose different experimental settings to simulate real-world scenarios. Part of the scientific leaderboards can be viewed as a special type of scientific knowledge graph that includes three types of entities (Task, Dataset, Metric) and the relations between them, which have been the primary focus of the previous studies on information extraction from scientific literature [5, 7, 8, 9, 10]. Our work in the cold start scenario, in which we do not assume any pre-defined TDM triple is given, constructs such a scientific knowledge graph and links the papers to the nodes in the graph simultaneously.

**EVALUATION:** <reasoning>The draft states the main paper's contribution and how it differs from existing literature. It outlines the limitations of previous studies and then explicitly states how the current work addresses these issues through specific contributions, such as manual dataset construction and handling cold start scenarios without pre-defined TDM triples. Therefore, the evaluation score should be 1.</reasoning><score>1</score>

<END OF EXAMPLE 1>

Figure 12: Prompt for Related Work Evaluation - Positioning Consistency (Each Paragraph) (Part 1)

### Prompt for Related Work Evaluation - Positioning Consistency (Each Paragraph) (Part 2)

<START OF EXAMPLE 2>

**ANSWER:** Several efforts have been made to automate the extraction and organization of performance results from scientific publications. Early work by [1] introduced a framework, TDMS-IE, for identifying task, dataset, metric, and score tuples from NLP papers, establishing a foundation for automatic leaderboard construction. Similarly, AxCell [2] proposes a robust pipeline that utilizes structural analysis and novel table segmentation techniques to extract results from machine learning papers, demonstrating significant improvements over prior methods. ORKG-Leaderboards [3] offers a systematic approach that integrates leaderboard extraction into a knowledge graph framework, enabling machine-actionable publishing and dynamic visualization of state-of-the-art performance. TELIN [4] focuses on extracting leaderboard-relevant entities from PDFs using a semi-automated approach that reduces human annotation needs through targeted entity refinement. In addition to these extraction systems, several datasets and benchmarks have been introduced to facilitate the development and evaluation of leaderboard construction tools. LEGOBench [5] provides a large-scale benchmark derived from arXiv and PapersWithCode, and evaluates both language model-based and graph-based approaches. SciERC and its associated framework SciIE [6] support multi-task extraction of entities and relations, enabling construction of scientific knowledge graphs. SciREX [7] extends information extraction to the document level, capturing relationships that span across sections, which is critical for leaderboard generation from full papers.

**EVALUATION:** <reasoning>The draft provides a comprehensive overview of related work but does not explicitly or implicitly mention the main paper’s own contribution or how it differs from the existing literature. It talks about the efforts made but doesn’t bridge into how this work builds upon them or addresses gaps they left. It focuses solely on summarizing prior work without establishing the unique position or advancement made by the current study. Therefore, the evaluation score should be 0.</reasoning>

<score>0</score>

<END OF EXAMPLE 2>

[ANSWER]: {Data instance to be evaluated}

Figure 13: Prompt for Related Work Evaluation - Positioning Consistency (Each Paragraph) (Part 2)

### Prompt for Related Work Evaluation - Positioning Consistency (Final Paragraph) (Part 1)

**[QUERY]:** Given a context related work paragraph, your task is to write a final paragraph of related work section for a scientific paper that states the paper’s contribution or its position among the literature.

**[CRITERIA]:** As a complementary to context paragraph, final paragraph should explicitly or implicitly state the current paper’s contributions or its position among the literature while addressing the points from the context paragraph. In other words, contributions in the final paragraph should include the discussed points in the context paragraph. Scoring rubric is as follows:

0: Final paragraph fails to address the points in context paragraph while stating the contributions and positioning or fails to state contributions and positioning at all.

1: Final paragraph successfully address the points in context paragraph while stating the contributions and positioning of the paper.

**[EXAMPLES]:**

<START OF EXAMPLE 1>

**CONTEXT:** In a parallel research thread, scientific leaderboards have been treated as a form of scientific knowledge graphs composed of Task, Dataset, and Metric entities. Information extraction in this context has been widely studied, as seen in works such as [5, 7, 8, 9, 10]. These studies typically focus on linking scientific content to structured entities but often assume that the schema or triples are known in advance.

**ANSWER:** In contrast to prior work, our approach addresses several critical limitations. We construct our dataset manually from publications to ensure comprehensive TDMR annotations, rather than relying on incomplete or unstandardized third-party sources. We normalize TDMR tuples for accurate leaderboard construction and propose novel experimental settings that reflect realistic, dynamic scenarios, including cold-start conditions without any pre-defined TDM taxonomy. By simultaneously constructing a scientific knowledge graph and linking publications to it in an open-domain setting, our method is inherently adaptable to the evolving research landscape.

**EVALUATION:** <reasoning>The context paragraph discusses prior work on scientific leaderboards as knowledge graphs with entities like Task, Dataset, and Metric (TDM). It highlights that prior studies focus on information extraction and linking to structured entities but usually assume a pre-existing schema or set of triples. The final paragraph contrasts the main paper’s contributions with this prior work by constructing a scientific knowledge graph and linking publications to it in an open-domain and adaptable manner, which directly relates to the context’s theme of schema-free extraction and linking. The final paragraph clearly positions the paper among existing literature and addresses the specific limitation mentioned in the context—i.e., the assumption of known schema—by proposing methods that work without one. Therefore, the evaluation score should be 1.</reasoning><score>1</score>

<END OF EXAMPLE 1>

Figure 14: Prompt for Related Work Evaluation - Positioning Consistency (Final Paragraph) (Part 1)

### Prompt for Related Work Evaluation - Positioning Consistency (Final Paragraph) (Part 2)

<START OF EXAMPLE 2>

**CONTEXT:** In a parallel research thread, scientific leaderboards have been treated as a form of scientific knowledge graphs composed of Task, Dataset, and Metric entities. Information extraction in this context has been widely studied, as seen in works such as [5, 7, 8, 9, 10]. These studies typically focus on linking scientific content to structured entities but often assume that the schema or triples are known in advance.

**ANSWER:** Building on existing lines of research, this work explores TDMR extraction and scientific leaderboard construction. The dataset is compiled from scientific publications. The approach is designed to accommodate variations in how benchmark results are reported.

**EVALUATION:** <reasoning>The context paragraph discusses prior research treating scientific leaderboards as knowledge graphs with a focus on information extraction where the schema is often assumed to be known. The main point raised is that existing methods typically presume a known structure for the extracted information. While the final paragraph mentions the extraction of TDMR and the construction of leaderboards—aligning with the topic in the context—it does not explicitly address whether this work assumes a known schema or introduces any method for schema discovery or flexibility. Thus, the final paragraph does not clearly state how it handles the assumption of a known schema, which is a key point of the context paragraph. Therefore, the evaluation score should be 0.</reasoning><score>0</score>

<END OF EXAMPLE 2>

[ANSWER]: {*Data instance to be evaluated*}

Figure 15: Prompt for Related Work Evaluation - Positioning Consistency (Final Paragraph) (Part 2)

### Prompt for Review Utility Evaluation - Verifiability Extraction

**[QUERY]:** Your task is to write a review comment for a scientific paper. The comment should contain claims, opinions, or suggestions that require justification.

**[CRITERIA]:** A statement is considered a claim if it includes any of the following: subjective statements, suggestions or requests for changes, judgments about sections of the paper, deductions or inferred observations, and any statement where evidence or justification is required to support the claim. A statement is considered a normal (non-claim) statement if it meets these criteria: It does not contain an opinion, claim, or suggestion but consists solely of factual, descriptive content, it indicates the existence or absence of something without suggesting changes, it makes general statements about the paper that do not express an opinion, it consists of objective, factual statements that do not require verification, it includes requests for clarification or general questions, it presents logical statements or directly inferable information. Scoring rubric is as follows:

0: If the text consists solely of normal statements.

1: If the text contains claims, opinions.

**[EXAMPLES]:**

<START OF EXAMPLE 1>

**ANSWER:** You mention that dataset X is used frequently in this domain, but there is no reference supporting this statement. Could you add one?

**EVALUATION:** <reasoning>The comment points out the lack of a reference and suggests adding one but does not assert a subjective claim requiring verification. Since it is a factual observation with a request, the evaluation score should be 0.</reasoning><score>0</score>

<END OF EXAMPLE 1>

<START OF EXAMPLE 2>

**ANSWER:** Lack of novelty:- Adversarial attacks by perturbing text has been done on many NLP models and image-text models. It is nicely summarized in related work of this paper. The only new effort is to take similar ideas and apply it on video-text models.

**EVALUATION:** <reasoning>The comment claims a "lack of novelty" regarding adversarial attacks, referencing the related work section. While it provides a logical basis, additional references would strengthen it, making the claim mostly verifiable. As it is a claim, its evaluation score should be 1.</reasoning><score>1</score>

<END OF EXAMPLE 2>

**[ANSWER]:** {*Data instance to be evaluated*}

Figure 16: Prompt for Review Utility Evaluation - Verifiability Extraction

## Prompt for Review Utility Evaluation - Verifiability (Part 1)

**[QUERY]:** Your task is to write a review comment for a scientific paper. The claims in your comment should be justified or proved by providing logical reasoning, using common sense, or referencing external sources.

**[CRITERIA]:** Claim justification-verification can be done either by logical reasoning supporting the claim, common sense knowledge in the field verifying the claim (e.g., referencing established practices or standards), or external references substantiating the claim. Scoring rubric is as follows:

- 1: The comment contains a claim without any supporting evidence or justification.
- 2: The comment provides some support for its claim, but the justification is vague, insufficient, or not fully articulated. Authors may struggle to follow the reasoning.
- 3: The comment provides support for its claim, but key elements are missing, such as specific examples, detailed explanations, or supporting references. Authors must make a significant effort to follow the justification.
- 4: The comment's claim is sufficiently supported but has minor gaps. The reviewer could provide a more detailed explanation or reference.
- 5: The claim is thoroughly supported by explicit, sufficient, and robust evidence. This can be achieved through: - Clear and precise reasoning or explanation. - Specific and relevant references to external works or data. - Logical and unassailable common-sense arguments.

**[EXAMPLES]:**

<START OF EXAMPLE 1>

**ANSWER:** A number of claims from this paper would benefit from more in-depth analysis.

**EVALUATION:** <reasoning>The comment states that "a number of claims from this paper would benefit from more in-depth analysis." However, it does not provide any supporting reasoning or examples to justify this opinion. As a result, the comment is unverifiable, since the lack of evidence and explanation makes it difficult for the authors to understand or address the concern. Therefore, the evaluation score should be 1.</reasoning><score>1</score>

<END OF EXAMPLE 1>

<START OF EXAMPLE 2>

**ANSWER:** The relatively poor performance on nouns makes me uneasy. While I can expect TWSI to do really well due to its nature, the fact that the oracle GAP for PPDBClus is higher than most clustering approaches is disconcerting, and I would like to understand the gap better. This also directly contradicts the claim that the clustering approach is generalizable to all parts of speech (124-126), since the performance clearly isn't uniform.

**EVALUATION:** <reasoning>The comment presents a claim that the clustering approach is not generalizable to all parts of speech due to non-uniform performance, which is meant to contradict the authors' claims. However, the explanation provided is minimal, stating only that the performance is "clearly isn't uniform." While this suggests an issue, the reasoning lacks sufficient detail and examples, making it hard for the authors to fully understand the justification and address it effectively. As a result, the claim is only borderline verifiable. Therefore, evaluation score should be 2.</reasoning><score>2</score>

<END OF EXAMPLE 2>

Figure 17: Prompt for Review Utility Evaluation - Verifiability (Part 1)

### Prompt for Review Utility Evaluation - Verifiability (Part 2)

<START OF EXAMPLE 3>

**ANSWER:** Lines 102-106 is misleading. While intersection and probs are true, "such distribution" cannot refer to the discussion in the above.

**EVALUATION:** <reasoning>The comment makes a claim that "Lines 102-106 are misleading" due to the use of "'such distribution'." While the reviewer begins to justify this by mentioning that it cannot refer to the prior discussion, the explanation is not sufficient for a full verification of the claim. More detailed reasoning or examples are needed to help the authors grasp the issue clearly, landing the comment at the evaluation score 3.</reasoning><score>3</score>

<END OF EXAMPLE 3>

<START OF EXAMPLE 4>

**ANSWER:** Also, since the dataset is artificially created, the dataset itself might have a lot of noise. For instance, the collected "pristine" set of tweets may not be pristine enough and might instead contain misinformation as well as out-of-context images. I would have liked to see more analysis around the quality of the collected dataset and the amount of noise it potentially has.

**EVALUATION:** <reasoning>The review comment makes a verifiable claim that the "dataset itself might have a lot of noise," citing the example that the 'pristine' set of tweets may include misinformation and out-of-context images. This reasoning provides a logical basis for the claim, making it mostly verifiable. However, while the argument is logically sound, it lacks detailed evidence or references, leaving minor gaps that could be addressed. Therefore, the evaluation score should be 4.</reasoning><score>4</score>

<END OF EXAMPLE 4>

<START OF EXAMPLE 5>

**ANSWER:** - line 47 - 48 "over-parametrization invariably overfits the data and results in worse performance": over-parameterization seems to be very helpful for supervised learning of deep neural networks in practice ... Also, I have seen a number of theoretical work showing the benefits of over-parametrisation e.g. [1].

**EVALUATION:** <reasoning>The review comment challenges the claim made in the paper that "over-parametrization invariably overfits the data and results in worse performance." The reviewer provides a counter-argument by stating that over-parameterization is often beneficial in supervised learning of deep neural networks and supports this with a reference to theoretical work that demonstrates these benefits. This provides a clear and precise reasoning, along with a reference, making the claim fully verifiable. Therefore, the comment is fully verifiable and should get the evaluation score 5.</reasoning><score>5</score>

<END OF EXAMPLE 5>

[ANSWER]: {Data instance to be evaluated}

Figure 18: Prompt for Review Utility Evaluation - Verifiability (Part 2)

### Prompt for Review Utility Evaluation - Actionability (Part 1)

**[QUERY]:** Your task is to write a review comment for a scientific paper. The comment should be actionable. Those actions should be clearly identifiable and concrete.

**[CRITERIA]:** Explicit actions or suggestions are direct or apparent. Authors can directly identify modifications they should apply to their draft. Clarification questions should be treated as explicit statements if they give a direct action. However, implicit actions need to be inferred from the comment. This includes missing parts that need to be added. Authors can deduce what needs to be done after reading the comment. For concrete actions, the authors know exactly what needs to be done and how to apply the action. However, for vague actions the authors still don't know how to carry out this action. Scoring rubric is as follows:

- 1: The comment lacks meaningful information to help authors improve the paper. Authors do not know what they should do after reading the comment.
- 2: The comment includes an implicitly stated action or an action that can be inferred. However, the action itself is vague and lacks detail on how to apply it.
- 3: The comment explicitly states an action but is vague on how to execute it.
- 4: The comment implicitly states an action but concretely states how to implement the inferred action.
- 5: The comment contains an explicit action and concrete details on how to implement it. Authors know exactly how to apply it.

**[EXAMPLES]:**

<START OF EXAMPLE 1>

**ANSWER:** The hGRU architecture seems pretty ad-hoc and not very well motivated.

**EVALUATION:** <reasoning>The review comment, "The hGRU architecture seems pretty ad-hoc and not very well motivated," lacks specificity and actionable guidance for the authors. While it expresses a concern about the hGRU architecture being "ad-hoc" and "not very well motivated," it does not provide any detailed explanation or examples of why the reviewer perceives it this way. Without specific points or suggestions, the authors are left without a clear understanding of what aspects of the architecture need further clarification or improvement. hence, this comment is not actionable at all. Therefore the evaluation score should be 1.</reasoning><score>1</score>

<END OF EXAMPLE 1>

<START OF EXAMPLE 2>

**ANSWER:** The paper does not dig into the theory proofs and show the convergence properties of the proposed algorithm.

**EVALUATION:** <reasoning>The comment highlights the need for the paper to delve deeper into the theory proofs and demonstrate the convergence properties of the proposed algorithm. However, it does not specify which aspects of the theory proofs require further exploration or how the authors should approach this task. The lack of concrete recommendations regarding the convergence properties leaves the authors uncertain about how to fulfill this request effectively. Due to its vague and indirect nature, this comment is barely actionable. Therefore the evaluation score should be 2.</reasoning><score>2</score>

<END OF EXAMPLE 2>

Figure 19: Prompt for Review Utility Evaluation - Actionability (Part 1)

### Prompt for Review Utility Evaluation - Actionability (Part 2)

<START OF EXAMPLE 3>

**ANSWER:** A number of claims from this paper would benefit from more in-depth analysis.

**EVALUATION:** <reasoning>The comment points out that certain claims require more in-depth analysis but does not clarify which claims need further scrutiny. As a result, the authors may not know where to focus their efforts, leading to potential misinterpretation of the feedback. Since the suggested action is direct but still lacks the necessary details for precise implementation, this comment is somewhat actionable. Therefore the evaluation score should be 3.</reasoning><score>3</score>

<END OF EXAMPLE 3>

<START OF EXAMPLE 4>

**ANSWER:** The adopted baseline models are weak. First of all, the author does not compare to Campos et al. (2020), which also uses feedback in QA tasks. Second, they also do no comparison with other domain adaptation methods, such as those work cited in Section 8.

**EVALUATION:** <reasoning>The reviewer identifies specific actions that the authors did not take, such as not comparing their work to Campos et al. (2020) and other domain adaptation methods cited in Section 8. This serves as an implicit suggestion that these comparisons should be implemented. However, rather than explicitly stating what the authors should do, the comment leaves room for interpretation. Because the required action is clear but not directly stated, the comment is mostly actionable. Therefore the evaluation score should be 4.</reasoning><score>4</score>

<END OF EXAMPLE 4>

<START OF EXAMPLE 5>

**ANSWER:** The abstract is written well and invokes intrigue early - could potentially be made even better if, for "evaluating with gold answers is inconsistent with human evaluation" - an example of the inconsistency, such as models get ranked differently is also given there.

**EVALUATION:** <reasoning>The comment explicitly states that an example of inconsistency should be provided in the abstract, specifically where it mentions "evaluating with gold answers is inconsistent with human evaluation." By directly instructing the authors to include an example, such as how models get ranked differently, it removes any uncertainty about how to proceed. Since the feedback is clear, specific, and directly actionable, the comment is fully actionable. Therefore the evaluation score should be 5.</reasoning><score>5</score>

<END OF EXAMPLE 5>

[ANSWER]: {Data instance to be evaluated}

Figure 20: Prompt for Review Utility Evaluation - Actionability (Part 2)

### Prompt for Review Utility Evaluation - Grounding-specificity (Part 1)

**[QUERY]:** Your task is to write a review comment for a scientific paper. The comment should refer to a specific part of the paper and clearly identify the issue with that part.

**[CRITERIA]:** For fully grounded comment, the author can accurately pinpoint the section, table, figure, or unique aspect being addressed. For weak grounded comment, the author can make an educated guess but cannot precisely identify the referenced part. For specificity, the comment should detail what is wrong or missing in the referenced part. If external work is mentioned, it should also provide specific examples. Scoring rubric is as follows:

- 1: The comment is not grounded at all. It does not identify a specific area in the paper. The comment is highly unspecific.
- 2: The authors cannot confidently determine which part the comment addresses. Further, the comment does not specify what needs to be addressed in this part.
- 3: The authors cannot confidently determine which part the comment addresses. However, the comment clearly specifies what needs to be addressed in this part.
- 4: The comment explicitly mentions which part of the paper it addresses, or it should be obvious to the authors. However, this comment does not specify what needs to be addressed in this part.
- 5: The comment explicitly mentions which part of the paper it addresses, and it is obvious to the authors. The comment specifies what needs to be addressed in this part.

**[EXAMPLES]:**

<START OF EXAMPLE 1>

**ANSWER:** The writing should be improved. Some points in the paper is unclear to me.

**EVALUATION:** <reasoning>The comment is not grounded as it does not identify any specific part of the paper that needs improvement. It generally mentions that the writing should be improved and that some points are unclear, but it does not specify which points or sections are problematic. This lack of detail makes it difficult for the authors to know which parts of the paper to focus on for revisions. Therefore, this comment is not grounded and highly unspecific, aligning with the evaluation score 1.</reasoning><score>1</score>

<END OF EXAMPLE 1>

<START OF EXAMPLE 2>

**ANSWER:** It is always easier to show something (i.e. attention in seq2seq MTL) is not working, but the value would lie in finding out why it fails and changing the attention mechanism so that it works.

**EVALUATION:** <reasoning>The authors have some idea about which part this comment addresses, as it refers to the attention mechanism in seq2seq MTL. However, they still do not know what specifically needs to be addressed in this part. While they can guess where they mentioned that something is not working, they still cannot identify the specific issues or suggestions for improvement. The comment suggests that the value lies in understanding why the attention mechanism fails and how it can be modified to work effectively. However, without specific guidance or examples, the authors are left without clear direction. Therefore, this comment is weakly grounded and not specific, aligning with the evaluation score 2.</reasoning><score>2</score>

<END OF EXAMPLE 2>

Figure 21: Prompt for Review Utility Evaluation - Grounding-specificity (Part 1)

### Prompt for Review Utility Evaluation - Grounding-specificity (Part 2)

<START OF EXAMPLE 3>

**ANSWER:** Would the use of feature engineering help in improving the performance? Uto et al. (2020)'s system reaches a QWK of 0.801 by using a set of hand-crafted features. Perhaps using Uto et al. (2020)'s same feature set could also improve the results of this work.

**EVALUATION:** <reasoning>The comment suggests a potential improvement in the system's performance through feature engineering, specifically referencing Uto et al. (2020) as an example. While it does not explicitly mention a specific section of the paper, the authors can infer that it relates to the discussion on performance. The suggestion is specific, as it points to a particular feature set that could be considered. Therefore, the comment is weakly grounded but specific, aligning with the evaluation score 3.</reasoning><score>3</score>

<END OF EXAMPLE 3>

<START OF EXAMPLE 4>

**ANSWER:** The paper's title might not accurately reflect the central theme and main contributions of your document. A title is crucial in drawing the interest of readers and offering them a clear understanding of the primary focus of your study, guiding them from the outset.

**EVALUATION:** <reasoning>The title is explicitly mentioned, grounding the comment. Despite this, it is under-specific because it doesn't indicate whether the problem is misleading wording, an inadequate reflection of scope, or the omission of critical terms that would better describe the paper's content. Hence, this comment is fully grounded and under-specific. Therefore the evaluation score should be 4.</reasoning><score>4</score>

<END OF EXAMPLE 4>

<START OF EXAMPLE 5>

**ANSWER:** Why the results of Table 6 is not aligned with Table 1 (MCT-pair)? Also what about the ablation studies of MCT without the adaptive metrics.

**EVALUATION:** <reasoning>The comment is fully grounded as it explicitly mentions "Table 6" and "Table 1," allowing the authors to accurately identify the parts of the paper being addressed. Additionally, it is specific because it details the issue of alignment between the results in these tables and inquires about the ablation studies of MCT without the adaptive metrics. Therefore, this comment is categorized as fully grounded and specific, aligning with the evaluation score of 5.</reasoning><score>5</score>

<END OF EXAMPLE 5>

[ANSWER]: {Data instance to be evaluated}

Figure 22: Prompt for Review Utility Evaluation - Grounding-specificity (Part 2)

### Prompt for Review Utility Evaluation - Helpfulness (Part 1)

**[QUERY]:** Your task is to write a review comment for a scientific paper. The comment should be useful for the authors to help improving the paper.

**[CRITERIA]:** A helpful review should be actionable, grounded on a specific part of the paper, provide justification or evidence to its claims. Scoring rubric is as follows:

- 1: The comment fails to identify meaningful weaknesses or suggest improvements, leaving the authors with no actionable feedback.
- 2: The comment identifies a weakness or improvement area but is vague, lacks clarity, or provides minimal guidance, making it only slightly beneficial for the authors.
- 3: The comment identifies weaknesses or areas for improvement but is incomplete or lacks depth. While the authors gain some insights, the feedback does not fully address their needs for improving the draft.
- 4: The comment provides clear and actionable feedback on weaknesses and areas for improvement, though it could be expanded or refined to be fully comprehensive and impactful.
- 5: The comment thoroughly identifies weaknesses and offers detailed, actionable, and constructive suggestions that empower the authors to significantly improve their draft.

**[EXAMPLES]:**

<START OF EXAMPLE 1>

**ANSWER:** Lack of novelty: Adversarial attacks by perturbing text has been done on many NLP models and image-text models. It is nicely summarized in related work of this paper. The only new effort is to take similar ideas and apply it on video-text models.

**EVALUATION:** <reasoning>This comment simply recaps a section already addressed in the paper and does not offer constructive guidance or specific suggestions on overcoming the stated issue of lack of novelty. Without actionable advice or an indication of how to differentiate or enhance the work, the comment is not useful to the authors and does not contribute to the paper's improvement, aligning with the evaluation score of 1: The comment is not helpful at all.</reasoning><score>1</score>

<END OF EXAMPLE 1>

<START OF EXAMPLE 2>

**ANSWER:** I don't think the probabilistic connection is drawn very well. It doesn't seem to be made formally enough to take it as anything more than motivational which is fine, but I would suggest the authors either cement this connection more formally or adjust the language to clarify.

**EVALUATION:** <reasoning>The comment highlights a potential issue with the way the probabilistic connection is presented in the paper, indicating it might lack formalism and appear purely motivational. While this could be valuable for the authors, the feedback is not specific enough to be fully actionable. The suggestion to either "cement this connection more formally" or "adjust the language to clarify" is a start, but it lacks specific guidance or examples that would make it more useful. Therefore, the authors are left with only a vague direction for improvement, making this feedback barely helpful. Therefore, the evaluation score should be 2.</reasoning><score>2</score>

<END OF EXAMPLE 2>

Figure 23: Prompt for Review Utility Evaluation - Helpfulness (Part 1)

## Prompt for Review Utility Evaluation - Helpfulness (Part 2)

<START OF EXAMPLE 3>

**ANSWER:** It is not clear if the proposed methodology is specific to bimanual manipulation. Just using robotic manipulation could be more appropriate.

**EVALUATION:** <reasoning>The comment identifies a potential issue with the specificity of the proposed methodology, questioning whether it is limited to bimanual manipulation or if it could be generalized to robotic manipulation. This is a relevant point that could help the authors clarify the scope and applicability of their methodology. However, the comment lacks depth and does not provide specific suggestions or guidance on how the authors might address this issue. It raises an important question but does not offer actionable steps for improvement, making it somewhat helpful. Therefore, the comment is evaluated as somewhat helpful, consistent with the evaluation score 3.</reasoning><score>3</score>

<END OF EXAMPLE 3>

<START OF EXAMPLE 4>

**ANSWER:** The authors need to perform ablation experiments to compare the proposed method with other methods (e.g., TubeR) in terms of the number of learnable parameters and GFLOPs.

**EVALUATION:** <reasoning>The review comment provides a clear and actionable suggestion for improvement by recommending the authors perform ablation experiments. It specifies the need to compare the proposed method with other methods, such as TubeR, focusing on the number of learnable parameters and GFLOPs. This feedback is direct and offers a concrete step for the authors to enhance their paper, making it mostly helpful. However, it could be further improved by explaining why these comparisons are important or how they might impact the paper's conclusions, which would elevate it to a highly helpful comment. Therefore, it aligns with a score of 4.</reasoning><score>4</score>

<END OF EXAMPLE 4>

<START OF EXAMPLE 5>

**ANSWER:** Finally, and similarly to above, i'd like to see an experiment where the image is occluded (half of the image is randomly blacked out). This (a) simulates the irregularity that is often present in neural/behavioral data (e.g. keypoint detection failed for some mice in some frames), and (b) would allow us to inspect the long-range "inference" capacity of the model, as opposed to a nearly-supervised reconstruction task. Again, these should be reasonably easy experiments to run. I'd expect to see all of these experiments included in a final version (unless the authors can convince me otherwise).

**EVALUATION:** <reasoning>The comment is highly helpful as it provides a clear and specific suggestion for an additional experiment that could significantly enhance the paper. By proposing an experiment where the image is occluded, the reviewer not only identifies a potential area for improvement but also explains the rationale behind it—simulating irregularities in data and testing the model's inference capacity. This feedback is actionable and constructive, offering the authors a concrete way to strengthen their work. The expectation that these experiments be included unless justified otherwise further emphasizes the importance of this suggestion, making it a valuable contribution to the paper's development.</reasoning><score>5</score>

<END OF EXAMPLE 5>

[ANSWER]: {Data instance to be evaluated}

Figure 24: Prompt for Review Utility Evaluation - Helpfulness (Part 2)

### Prompt for Novelty Alignment Evaluation (Part 1)

**[QUERY]:** Your task is to write two assessments regarding novelty of a scientific paper. Each one should lead to implicitly/explicitly a verdict of either "novel" or "not novel". The assessments should be aligned in terms of their novelty decision.

**[CRITERIA]:** Assessments with verdict "novel" state that the paper introduces new, original, or significant ideas, methods, metrics, or frameworks; emphasizes a meaningful, notable, or important contribution; describes the work as advancing the field in a substantive or distinct way. Even if it is not groundbreaking fundamental new contribution, the paper can be still novel based on its significant contributions. On the other hand, assessments implying "not novel" state that the contribution is totally incremental, limited, weak, already known; emphasize that prior work already covers most of the ideas or findings; say the paper lacks significant, substantial, or original contributions; describe the work as mainly empirical, confirmatory, or replicating existing knowledge without new insights. You need to check whether these two assessments' final conclusion is the same or not. It is possible that two assessments can have similar observations on some points but they can arrive different conclusions. Your evaluation should be according to final conclusions. Evaluation two assessments should be binary. Scoring rubric is as follows.

0: Two assessments do not come to the same conclusion in terms of novelty verdict.

1: Two assessments come to the same conclusion in terms of novelty verdict.

**[EXAMPLES]:**

<START OF EXAMPLE 1>

**ANSWER: Assessment 1:** This paper introduces a novel framework for analyzing the self-improvement capabilities of large language models, centering on the proposal of the generation-verification gap (GV-gap) as a new metric to quantify the limits of self-improvement. The GV-gap is clearly defined and illustrated with real-world examples, offering a fresh perspective for measuring and understanding where self-improvement may be fundamentally constrained. This approach represents a meaningful contribution to the research direction, as it provides the community with a new tool for future studies on model self-improvement. While the concept of self-improvement in language models has been explored in prior work, the introduction of the GV-gap offers a distinct and unified way to formalize and assess these capabilities. The novelty of the paper lies in this new quantification method, which can help clarify and advance discussions in the field. However, the practical application of the GV-gap is somewhat limited by the noisiness of real-world utility functions, as acknowledged by the authors, which may affect the robustness of the measurements. Despite this limitation, the proposed framework and metric are likely to be valuable for future research, marking the work as a notable and useful contribution to ongoing efforts in understanding and improving language model self-improvement. **Assessment 2:** This submission offers a comprehensive and systematic empirical/theoretical study of LLM self-improvement, with its main novelty being the formalization and central use of the "generation-verification gap" (GV-Gap) as a unifying metric. While the concept of a gap between generation and verification is present in prior work, the explicit metric and its application across models and tasks, as well as the discovery of a scaling law for GV-Gap, are new contributions. The paper also provides a detailed, cross-model analysis of verification mechanisms, including ensemble verification, which is a substantive but incremental extension of existing meta-judging and reward aggregation methods. However, the submission tends to overstate the lack of systematic analysis and diversity in prior work, and does not fully engage with risks such as bias and diversity collapse highlighted in the literature. Overall, the work represents a significant incremental advance in empirical rigor and formalization, rather than a fundamentally new paradigm for LLM self-improvement.

Figure 25: Prompt for Novelty Alignment Evaluation (Part 1)

### Prompt for Novelty Alignment Evaluation (Part 2)

**EVALUATION:** <reasoning>Assessment 1 clearly concludes that the work is novel, emphasizing that the GV-gap constitutes a "meaningful contribution" and "new metric," marking the paper as a "notable and useful contribution." Assessment 2, while calling the work "a significant incremental advance," still attributes new contributions (formalizing GV-Gap, discovering a scaling law). It does not declare the work non-novel or merely confirmatory; it frames the contribution as incremental "but still novel". Thus both assessments align on a "novel" verdict, the score should be 1.</reasoning><score>1</score>

<END OF EXAMPLE 1>

<START OF EXAMPLE 2>

**ANSWER: Assessment 1:** This paper presents a comprehensive experimental analysis of self-improvement in Large Language Models (LLMs), focusing on the concept of the generation-verification gap and its relationship to model pre-training computational effort. While the study offers a modular framework and conducts extensive experiments to examine scaling phenomena and conditions for self-improvement, its novelty is limited. Most of the conclusions, such as the monotonic scaling of the verification gap with pre-training FLOPs and the identification of saturation limits, are already established in the literature. As a result, the paper does not provide fundamentally new insights or advances beyond what is already known. The contribution is primarily empirical, and while the analysis is thorough, it does not introduce novel theoretical perspectives or experimental findings that significantly advance the field. Consequently, the work's impact on the community is relatively weak from a novelty standpoint. **Assessment 2:** This submission offers a comprehensive and systematic empirical/theoretical study of LLM self-improvement, with its main novelty being the formalization and central use of the "generation-verification gap" (GV-Gap) as a unifying metric. While the concept of a gap between generation and verification is present in prior work, the explicit metric and its application across models and tasks, as well as the discovery of a scaling law for GV-Gap, are new contributions. The paper also provides a detailed, cross-model analysis of verification mechanisms, including ensemble verification, which is a substantive but incremental extension of existing meta-judging and reward aggregation methods. However, the submission tends to overstate the lack of systematic analysis and diversity in prior work, and does not fully engage with risks such as bias and diversity collapse highlighted in the literature. Overall, the work represents a significant incremental advance in empirical rigor and formalization, rather than a fundamentally new paradigm for LLM self-improvement.

**EVALUATION:** <reasoning>Assessment 1 concludes the paper is not novel, emphasizing that it offers only empirical, already-known findings, and no fundamentally new insights. Assessment 2 acknowledges the work as a **\*\*significant incremental advance\*\*** with **\*new contributions\*** (formalizing the metric, discovering a scaling law). Although incremental, it is still treated as novel. Since assessments do not align in terms of novelty of the paper, the evaluation score should 0.</reasoning><score>0</score>

<END OF EXAMPLE 2>

[ANSWER]: **Assessment 1:** {*First novelty assessment*} **Assessment 2:** {*Second novelty assessment*}

Figure 26: Prompt for Novelty Alignment Evaluation (Part 2)

### Prompt for Revision Evaluation - Relatedness (Part 1)

**[QUERY]:** Your task is to revise a scientific text according to given an instruction. The revised text should address the points in the instruction.

**[CRITERIA]:** The generated revision should correctly address the instruction. Scoring rubric is as follows.

0: The model revision does not address the instruction.

1: The revision follows the requirement of the instruction.

**[EXAMPLES]:**

<START OF EXAMPLE 1>

**ORIGINAL TEXT:** MNIST dataset with 1000 training images. All external regularization schemes except learning rate decay and batch normalization have been turned off. We perform the following experiments : 1 ) Full-batch gradient descent with  $\beta = 0$  (i.e., GD) for various learning rate  $h$  and the best test accuracy is noted (in Figure 2) to be 95 . Full-batch gradient descent with momentum (GD+M) performed for various  $\beta$  with a fixed step-size  $h = 0.1$  and the best test-accuracy is noted (in Figure 3) to be 96 . Our observation is that the best performance of GD (across all learning rates) is worse than the best performance of (GD+M) (across all  $\beta$  's). (Cohen et al., 2021) showed that gradient descent (GD) has an overwhelming tendency to increase the sharpness  $\lambda$  till it reaches  $2/h$ , called “the edge of stability”. And for (GD+M), the sharpness can reach up to  $2/(1+\beta)h$ , hence allowing it to enter a sharper region before becoming unstable. As greater allowable sharpness for (GD+M) than that of (GD) may suggest a higher test accuracy for (GD), this is not what we observe from the above experiment. We think the implicit regularization for (GD+M) plays a part in it. We believe IGR for momentum outweighs the sharpness effect in achieving better test accuracy.

**INSTRUCTION:** Summarize the second half of the paragraph to make the paragraph shorter.

**ANSWER:** MNIST dataset with 1000 training images. All external regularization schemes except learning rate decay and batch normalization have been turned off. We perform the following experiments : 1 ) Full-batch gradient descent with  $\beta = 0$  (i.e., GD) for various learning rates and the best test accuracy is noted to be 95. For (GD+M) performed for various  $\beta$  with a fixed step-size  $h = 0.1$  and the best test-accuracy is noted to be 96. Our observation is that the best performance of GD (across all learning rates) is worse than the best performance of (GD+M) across all  $\beta$  's).

**EVALUATION:** <reasoning>The provided answer omits the entire theoretical discussion rather than summarizing it, so it does not properly address the instruction. Therefore the evaluation score should be 0.</reasoning><score>0</score>

<END OF EXAMPLE 1>

Figure 27: Prompt for Revision Evaluation - Relatedness (Part 1)

### Prompt for Revision Evaluation - Relatedness (Part 2)

<START OF EXAMPLE 2>

**ORIGINAL TEXT:** MNIST dataset with 1000 training images. All external regularization schemes except learning rate decay and batch normalization have been turned off. We perform the following experiments : 1 ) Full-batch gradient descent with  $\beta = 0$  (i.e., GD) for various learning rate  $h$  and the best test accuracy is noted (in Figure 2) to be 95 . Full-batch gradient descent with momentum (GD+M) performed for various  $\beta$  with a fixed step-size  $h = 0.1$  and the best test-accuracy is noted (in Figure 3) to be 96 . Our observation is that the best performance of GD (across all learning rates) is worse than the best performance of (GD+M) (across all  $\beta$  's). (Cohen et al., 2021) showed that gradient descent (GD) has an overwhelming tendency to increase the sharpness  $\lambda$  till it reaches  $2/h$ , called “the edge of stability”. And for (GD+M), the sharpness can reach up to  $2(1 + \beta)/h$ , hence allowing it to enter a sharper region before becoming unstable. As greater allowable sharpness for (GD+M) than that of (GD) may suggest a higher test accuracy for (GD), this is not what we observe from the above experiment. We think the implicit regularization for (GD+M) plays a part in it. We believe IGR for momentum outweighs the sharpness effect in achieving better test accuracy.

**INSTRUCTION:** Summarize the second half of the paragraph to make the paragraph shorter.

**ANSWER:** MNIST dataset with 1000 training images. All external regularization schemes except learning rate decay and batch normalization have been turned off. We perform experiments with full-batch gradient descent (GD) and full-batch gradient descent with momentum (GD+M) for various learning rates and  $\beta$  values. Our results show that GD+M outperforms GD, achieving a best test accuracy of 96 compared to 95 for GD. This contradicts the expectation that GD+M's higher allowable sharpness would lead to better test accuracy. We attribute this to the implicit regularization effect of momentum, which outweighs the sharpness effect in achieving better test accuracy.

**EVALUATION:** <reasoning>The revision clearly summarizes the theoretical discussion in a shorter form while preserving the key ideas: GD+M's higher accuracy, the contradiction with sharpness-based expectations, and the role of implicit regularization. Because it successfully condenses the second half of the paragraph as instructed, it meets the requirement. Therefore the evaluation score should be 1.</reasoning><score>1</score>

<END OF EXAMPLE 2>

[**ORIGINAL TEXT**]: {Original text to be revised}

[**INSTRUCTION**]: {Revision instruction}

[**ANSWER**]: {Revision to be evaluated}

Figure 28: Prompt for Revision Evaluation - Relatedness (Part 2)

### Prompt for Revision Evaluation - Correctness (Part 1)

**[QUERY]:** Your task is to revise a scientific text according to given an instruction. The revised text should be in a better state after applying the instructions.

**[CRITERIA]:** The revised text should be a good quality revision proposition that should replace the original paragraph in the scientific article. Scoring rubric is as follows.

0: The revision is not a better version that can replace the original paragraph in the scientific article.

1: The revision is a better version that should replace the original paragraph in the scientific article.

**[EXAMPLES]:**

<START OF EXAMPLE 1>

**ORIGINAL TEXT:** SuperWeights, which are linear combination of templates which get reused by multiple layers . These SuperWeights capture a single operation on the input features ( e.g ., edge or texture detectors), and are themselves generated via a weighted combination of one or more templates of trainable parameters held by Weight Templates. Thus, to generate the weights for a single layer, we must first construct SuperWeights from the trainable parameters held by Weight Templates (discussed in Section 3.1), and then concatenate together all SuperWeights used by the layer to create its final weights (process illustrated in center-right column of Figure 2).

**INSTRUCTION:** Make the first half of the paragraph shorter by merging the two sentences and removing the details about how superweights are generated.

**ANSWER:** SuperWeights, linear combinations of templates reused by multiple layers, capture a single operation on input features (e.g., edge or texture detectors). Generated via weighted combinations of trainable parameters held by Weight Templates, they are then concatenated to create a layer’s final weights.

**EVALUATION:** <reasoning>The instruction asks to make the first half shorter by merging the two sentences and \*removing the details about how SuperWeights are generated\*. The provided answer does merge the sentences and shortens the text, but it still includes the generation detail (“Generated via weighted combinations of trainable parameters held by Weight Templates”), which the instruction explicitly requests to remove. Because the revision does not fully follow the instruction and retains disallowed details, it is not an improved replacement for the scientific article. Therefore, the evaluation score should be 0.</reasoning><score>0</score>

<END OF EXAMPLE 1>

Figure 29: Prompt for Revision Evaluation - Correctness (Part 1)

### Prompt for Revision Evaluation - Correctness (Part 2)

<START OF EXAMPLE 2>

**ORIGINAL TEXT:** SuperWeights, which are linear combination of templates which get reused by multiple layers . These SuperWeights capture a single operation on the input features ( e.g ., edge or texture detectors), and are themselves generated via a weighted combination of one or more templates of trainable parameters held by Weight Templates. Thus, to generate the weights for a single layer, we must first construct SuperWeights from the trainable parameters held by Weight Templates (discussed in Section 3.1), and then concatenate together all SuperWeights used by the layer to create its final weights (process illustrated in center-right column of Figure 2).

**INSTRUCTION:** Make the first half of the paragraph shorter by merging the two sentences and removing the details about how superweights are generated.

**ANSWER:** SuperWeights, which are linear combinations of templates, capture a single operation on the input features (e.g., edge or texture detectors). Thus, to generate the weights for a single layer, we must first construct SuperWeights from the trainable parameters held by Weight Templates (discussed in Section 3.1), and then concatenate together all SuperWeights used by the layer to create its final weights (process illustrated in center-right column of Figure 2).

**EVALUATION:** <reasoning>The instruction requires merging the first two sentences and removing details about how SuperWeights are generated. The revision successfully merges the sentences into one concise sentence describing what SuperWeights are and what they capture. It also removes the explicit description of their generation process. The remainder of the paragraph is kept intact and coherent. Overall, the revision is cleaner, follows the instruction, and is suitable to replace the original scientific text. Therefore, the evaluation score should be 1.</reasoning><score>1</score>

<END OF EXAMPLE 2>

[**ORIGINAL TEXT**]: {Original text to be revised}

[**INSTRUCTION**]: {Revision instruction}

[**ANSWER**]: {Revision to be evaluated}

Figure 30: Prompt for Revision Evaluation - Correctness (Part 2)