RUBRICS AS REWARDS: REINFORCEMENT LEARNING BEYOND VERIFIABLE DOMAINS

Anonymous authorsPaper under double-blind review

ABSTRACT

Reinforcement Learning with Verifiable Rewards (RLVR) has proven effective for complex reasoning tasks with clear correctness signals such as math and coding. However, extending it to real-world reasoning tasks is challenging, as evaluation depends on nuanced, multi-criteria judgments rather than binary correctness. Instance-specific rubrics have recently been used in evaluation benchmarks to capture such judgments, but their potential as reward signals for on-policy posttraining remains underexplored. We introduce **Rubrics as Rewards** (*RaR*), an on-policy reinforcement learning method that extends RLVR beyond verifiable domains by using rubric-based feedback. Across both medical and science domains, we evaluate multiple strategies for aggregating rubric feedback into rewards. The best RaR variant achieves relative improvements of up to 31% on HealthBench and 7% on GPQA-Diamond over popular LLM-as-judge baselines that rely on direct Likert-based rewards. These results demonstrate that RaR-trained policies adapt well to diverse evaluation formats, performing strongly on both rubric-based and multiple-choice tasks. Moreover, we find that using rubrics as structured reward signals yields better alignment for smaller judges and reduces performance variance across judge scales.

1 Introduction

Reinforcement Learning with Verifiable Rewards (RLVR) has enabled large language models to elicit complex reasoning on tasks with clear verifiable outcomes. This is especially effective in domains like math and code, where reward models can be replaced by scoring functions or test cases that automatically verify correctness (Lambert et al., 2024; Guo et al., 2025a; Cui et al., 2025). However, extending RLVR to unstructured, real-world reasoning is challenging because such tasks lack easily verifiable answers. A common workaround is to use preference-based reward models, but they tend to overfit superficial artifacts (e.g. response length, formatting quirks, annotator biases) (Singhal et al., 2023; Wang et al., 2024; Chen et al., 2024b; Ye et al., 2024; Gudibande et al., 2023) and require large volumes of pairwise comparisons (Ouyang et al., 2022). Instance-specific rubrics have recently emerged for nuanced evaluation in expert domains (Arora et al., 2025), yet their application in on-policy training for expert-level reasoning is largely unexplored.

To address this gap, we explore a paradigm shift that introduces a middle ground between the simplicity of verifiable rewards and the expressiveness of preference rankings, which often come with human artifacts and operational overhead. We introduce **Rubrics as Rewards (RaR)**, a framework for on-policy Reinforcement Learning that uses structured criteria or *rubrics* as the core reward mechanism. Rather than using rubrics only for evaluation (Arora et al., 2025; Sirdeshmukh et al., 2025), we treat them as checklist-style supervision that produces reward signals for on-policy RL. Each rubric is composed of modular, interpretable subgoals that provide automatable feedback aligned with expert intent. By decomposing "what makes a good response" into tangible, human-interpretable criteria, rubrics offer a middle ground between binary correctness signals and coarse preference rankings.

Previous works train generative reward models that learn to evaluate reasoning or final outputs with interpretable scores (Chen et al., 2025; Whitehouse et al., 2025; Anugraha et al., 2025; Guo et al., 2025b), and some have even used a model's internal confidence estimates as a proxy for reward (Zhao et al., 2025). More recent efforts have extended verifiable datasets beyond STEM domains,

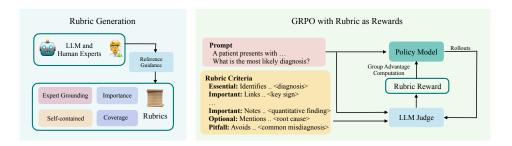


Figure 1: Overview of Rubrics as Rewards (RaR). (i) Rubric Generation: We synthesize prompt-specific, self-contained rubric criteria using a strong LLM guided by four core design principles, with reference answers serving as proxies for expert supervision. (ii) GRPO Training: These rubrics are used to prompt an LLM judge for reward estimation, which drives policy optimization via the GRPO on-policy learning loop.

broadening the applicability of RLVR methods to a wider range of tasks (Su et al., 2025b; Ma et al., 2025). Yet a general-purpose approach for specifying reliable reward signals remains elusive, particularly in tasks without a single ground truth where both subjective and objective criteria must be considered. In contrast, we treat rubrics as instance-specific, reusable reward functions. Once generated, rubrics provide interpretable and automatable supervision that can be applied consistently across new rollouts, offering a scalable and transparent alternative to opaque reward modeling in onpolicy learning.

Recent concurrent works explore checklists and principled rubric criteria for preference tuning and LLM safety (Gallego, 2025; Viswanathan et al., 2025; Dineen et al., 2025), highlighting a growing trend toward structured supervision. In contrast, we convert rubrics into reward functions for onpolicy RL, targeting expert reasoning and applied real-world domains. This closes the rubric-to-learning loop and improves performance on both rubric-guided evaluations and tasks with verifiable answers. Figure 1 illustrates our framework.

Our key contributions are as follows: (i) We introduce **Rubrics as Rewards** (**RaR**), an on-policy reinforcement learning framework that uses checklist-style rubrics for multi-criteria supervision in reasoning and real-world domains. (ii) We synthesize instance-specific rubrics for medicine and science and release the corresponding training sets, *RaR-Medicine* and *RaR-Science*. (iii) *RaR*-trained models consistently outperform strong baselines and yield a stable, generalizable training signal, with gains on both rubric-scored and verifiable multiple-choice evaluation settings. (iv) Our results demonstrate that rubric-based rewards provide stable supervision across judge sizes, helping smaller models align effectively with human preferences and maintaining robust evaluation performance from small to large judges.

2 Rubrics as Rewards

2.1 PROBLEM FORMULATION

Let x denote an input prompt and $\hat{y} \sim \pi_{\theta}(\cdot \mid x)$ be a sampled response from a model parameterized by θ . In domains without single ground-truth answers or automatic correctness signals, we define a structured reward function using instance-specific rubric criteria.

Each prompt x is associated with a set of k rubric items $\{(w_j, c_j)\}_{j=1}^k$, where $w_j \in \mathbb{R}$ denotes the weight of criterion j, and $c_j : (x, \hat{y}) \mapsto \{0, 1\}$ is a binary correctness function that indicates whether the response \hat{y} satisfies that criterion given the prompt.

2.2 REWARD AGGREGATION STRATEGIES

We investigate two complementary approaches for combining rubric feedback into scalar rewards:

Explicit Aggregation. Each criterion is independently evaluated using an LLM-as-judge, and the final normalized reward is computed as:

 $r(x,\hat{y}) = \frac{\sum_{j=1}^{k} w_j \cdot c_j(x,\hat{y})}{\sum_{j=1}^{k} w_j}$ (1)

Normalization makes rewards comparable across prompts that differ in rubric count or weights. Although we use binary checks for c_j in our experiments, the formulation can be extended to continuous-valued scores.

Implicit Aggregation. All rubric criteria along with categorical weights are passed to an LLM-as-judge, delegating the aggregation to the model itself to produce a single scalar reward:

$$r_{\text{implicit}}(x, \hat{y}) = f_{\phi}(x, \hat{y}, \{d_j\}_{j=1}^k)$$

$$\tag{2}$$

Here, f_{ϕ} denotes an LLM-based judge that takes the prompt x, the response \hat{y} , and the set of rubric criteria $\{d_j\}$ as input. This formulation allows the model to compute a holistic reward score directly, avoiding the need to manually tune rubric weights.

The prompts used for each method are detailed in Appendix A.7.

2.3 GENERALIZATION OF RLVR WITH RUBRICS AS REWARDS

Rubric-based reinforcement learning extends the standard RLVR (Reinforcement Learning with Verifiable Rewards) setting by supporting multi-dimensional, prompt-specific evaluation criteria. We formalize this relationship below.

Remark 1 (**Rubrics as Rewards subsumes RLVR**). The RLVR setting is a special case of rubricbased rewards defined in Equation 1, where k = 1, $w_1 = 1$, and $c_1(x, \hat{y})$ reduces to a single verifiable correctness function that compares the model output \hat{y} against the known correct answer y. For example, this could involve exact match or test case execution. Formally:

$$r_{\text{RLVR}}(x, \hat{y}) = \text{match}(y, \hat{y})$$
 (3)

where $match(y, \hat{y}) \in \{0, 1\}$ indicates whether the response satisfies the verifiable correctness condition.

Rubric-based reward functions thus generalize RLVR by enabling multi-dimensional supervision, flexible weighting across criteria, and the incorporation of both objective and subjective aspects of response quality. This formalization highlights that RLVR can be seen as a restricted instance of rubric-guided RL with a single essential criterion. In contrast, rubric-based rewards further enable structured supervision in settings where correctness is multifaceted and may not be strictly verifiable.

3 RUBRIC GENERATION

3.1 DESIDERATA

A rubric specifies criteria for high-quality responses and provides human-interpretable supervision. We identify four desiderata for effective rubric generation:

Grounded in Expert Guidance. Rubrics should reflect domain expertise by capturing the essential facts, reasoning steps, and conclusions necessary for correctness. Ideally, this grounding comes from human experts or their high-quality proxies.

Comprehensive Coverage. Rubrics should span multiple dimensions of response quality, including factual accuracy, logical coherence, completeness, style, and safety. Negative criteria (*pitfalls*) help identify frequent or high-risk errors that undermine overall quality.

Criterion Importance. Rubrics should reflect that some dimensions of response quality are more critical than others. For example, factual correctness must outweigh secondary aspects such as stylistic clarity. Assigning weights to criteria ensures this prioritization, whether through simple categorical tags, explicit numeric values, or learned weighting schemes.

Self-Contained Evaluation. Each rubric item should be independently actionable, allowing either human annotators or automated judges to assess it in isolation without requiring external context or domain-specific knowledge.

3.2 RUBRICS CREATION

We apply these desiderata to datasets for reasoning tasks in medicine and science. Given the scarcity of human-annotated rubric datasets in these domains, we use LLMs to generate instance-specific rubrics from golden reference answers at scale, enabling the study of structured rewards without costly human annotation.

For each prompt, an LLM generates a rubric of 7–20 self-contained items. Each item is assigned both a numeric and a categorical weight reflecting its relative importance. While numeric weights provide fine-grained prioritization, in our experiments we adopt categorical labels (*Essential*, *Important*, *Optional*, *Pitfall*) for ease of implementation and interpretability in controlled settings. The resulting rubrics are then used directly as reward functions through either explicit aggregation (Eq. 1) or implicit aggregation (Sec. 2.2).

In practice, we generate rubrics using OpenAI's o3-mini and GPT-40 (OpenAI o3-mini, 2025; Jaech et al., 2024; Hurst et al., 2024), conditioning generation on reference answers from the underlying datasets to approximate expert grounding. The resulting collections—*RaR-Medicine* and *RaR-Science*—are released for public use. These rubric sets supervise smaller policies under GRPO using both explicit and implicit reward aggregation.

4 EXPERIMENTS

4.1 Datasets

We investigate the utility of rubrics as rewards across two reasoning domains, medicine and science.

- RaR-Medicine: A dataset of 20k prompts drawn from diverse medical reasoning sources, including medical-o1-reasoning-SFT (Chen et al., 2024a), natural_reasoning (Yuan et al., 2025), SCP-116K (Lu et al., 2025), and GeneralThought-430K (General Reasoning, 2025). Instance-specific rubrics for this dataset are generated with GPT-40 (see Appendix A.2).
- RaR-Science: A dataset of ~20k prompts curated to align with GPQA-Diamond categories. Prompts are sourced from natural_reasoning (Yuan et al., 2025), SCP-116K (Lu et al., 2025), and GeneralThought-430K (General Reasoning, 2025), covering a broad range of scientific reasoning tasks (Appendix A.3). Rubrics for this dataset are synthesized with o3-mini.

4.2 Training Details

We conduct all experiments using on-policy reinforcement learning with the GRPO algorithm (Shao et al., 2024), taking Qwen2.5-7B as the base policy. Models are trained with a batch size of 96, a learning rate of 5×10^{-6} , and a constant schedule with 10% linear warmup. Complete hyperparameter settings are listed in Appendix A.4. Training runs are executed on a single compute node equipped with 8 NVIDIA H100 GPUs.

Our training pipeline consists of the following key components:

Response Generation: For each prompt q, we sample k=16 responses from the current policy π_{θ} , using a context length of 3584 and a sampling temperature of 1.0.

Reward Computation with Rubrics: We use gpt-4o-mini as the judge model to assign rewards R_q to the sampled responses. We experiment with various reward computation and aggregations strategies further described in Sections 4.3 and 4.4.

Policy Update: The policy weights are updated using GRPO based on the computed rewards.

¹Datasets will be made publicly available.

4.3 RUBRIC-FREE BASELINES

We consider various rubric-free baselines and off-the-shelf post-trained models. Rubric-free baselines are trained with Qwen2.5-7B as the base policy.

OFF-THE-SHELF: For off-the-shelf baselines we evaluate performance on Qwen2.5-7B. We also include the performance of Qwen2.5-7B-Instruct to compare with instruction-tuned variant of the base policy.

DIRECT-LIKERT: An LLM-as-judge provides a direct assessment for each response–prompt pair on a 1-10 Likert scale (Zheng et al., 2023; Kim et al., 2024), normalized to [0,1]. The resulting score is used directly as the reward signal for training.

REFERENCE-LIKERT: An LLM-as-judge compares the generated response against a reference answer (written by experts or stronger LLMs) and assigns a 1-10 Likert score (Zheng et al., 2023), normalized to [0,1]. This reference-guided score is used as the reward signal for policy updates. The reward for each (*prompt, response, reference*) triplet is defined as:

$$R_{\text{ref}}(q, x) = \text{Norm}(\text{LikertScore}(q, x, x^*))$$

where x^* denotes the reference answer.

4.4 Rubric-Guided Methods

RaR-PREDEFINED: This method uses a fixed set of generic rubrics for all prompts (e.g. *response is concise, response contains correct information*). It employs the Explicit Aggregation method (Equation 1) with all criteria weighted uniformly (see Appendix A.6).

RAR-EXPLICIT: This variant also uses Explicit Aggregation using a weighted sum (Equation 1) but applies it to instance-specific rubrics from Section 3. We manually assign numerical weights based on the generated categorical labels: {"Essential": 1.0, "Important": 0.7, "Optional": 0.3, "Pitfall": 0.9}².

RAR-IMPLICIT: This variant uses the Implicit Aggregation method (Equation 2). It leverages prompt-specific rubrics, where a judge model evaluates the response as a whole to assign a single Likert rating (1-10), avoiding the need for hand-tuned weights. The reward is normalized to the [0,1] range during training.

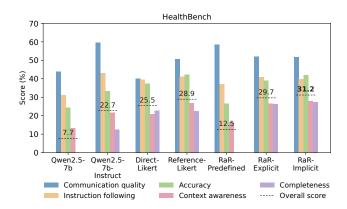
4.5 EVALUATION SETUP

Rubric-Based Evaluation We evaluate models trained with *RaR-Medicine* on Health-Bench (Arora et al., 2025), a benchmark of 5,000 clinical conversations designed to assess model safety and helpfulness in realistic medical scenarios. Performance is measured using detailed, physician-authored rubrics. We generate responses with greedy decoding (temperature = 0) and report both overall scores and per-axis scores following the original setup. For ablation studies, we sample a subset of 1,000 prompts (hereafter referred as HealthBench-1k) and use the rest for training.

Multiple-Choice Evaluation Each model is evaluated across 10 independent runs, using greedy decoding (*temperature=0*) to sample one response per prompt. Answer choices are permuted per example to reduce positional bias, and outputs are parsed for boxed answer formats (e.g., boxed{A}). If extraction fails, we fall back to a GPT-40 verifier that checks whether the response contains the correct option letter or text (see Appendix A.5). Final accuracy is reported as the mean over 10 runs, and we include 95% confidence intervals to account for run-to-run variance.

LLM-Judge Alignment Evaluation To measure how well LLM judges align with human preferences, we build a paired evaluation set from roughly 3,000 HealthBench prompts. For each prompt, we take the practitioner-approved answer as the *preferred* response and create a *perturbed* alternative via controlled edits (see Appendix A.10 for method used for perturbation and prompt selection). The metric is *pairwise preference accuracy* i.e. the fraction of pairs where the preferred response scores higher reported across judge models of varying sizes.

²Pitfall criteria are phrased in positive form (e.g., "The response avoids misinformation"), so satisfying them contributes positively to the score. If a pitfall is not satisfied, the corresponding reward is reduced or penalized.



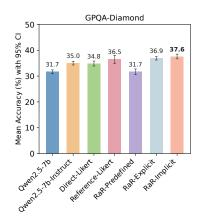


Figure 2: Performance of baselines and RaR (*Rubrics as Rewards*) variants for the medicine and science domains. **HealthBench (left):** shows per-axis scores across five core axes, with a thin dashed gray line indicating the overall score (all values shown as percentages). **GPQA-Diamond (right):** mean accuracy over 10 runs; error bars represent 95% confidence intervals. All policies are evaluated using gpt-4o-mini as the LLM-as-Judge. Across both domains, RaR-Implicit consistently outperforms Direct-Likert and demonstrates a competitive advantage over Reference-Likert.

5 RESULTS

We now present the main findings of our study.

Rubrics as Rewards shows strong gains across evaluation settings. Table 2 reports results on HealthBench (rubric-based, free-form) and GPQA-Diamond (multiple-choice). RaR-Implicit consistently outperforms Direct-Likert, with relative gains up to 31% on HealthBench and 7% on GPQA. Both rubric-guided variants achieve higher scores than the base and instruction-tuned policies. Gains on GPQA-Diamond show that rubric-induced skills generalize beyond rubric-based evaluation. The RaR-Predefined variant, which applies a fixed list of generic rubrics to every prompt (no instance-specific synthesis), underperforms because generic criteria miss prompt-specific requirements and common failure modes, producing misaligned reward signals. Hence, effective training requires instance-specific rubric synthesis as they better capture task context and typical failure modes.

Beyond these gains, RaR-Implicit also shows small but consistent gains over Reference-Likert. In our setup, rubrics are generated with stronger LLMs using reference answers as proxies for expert supervision, so rubric quality is impacted by reference quality. Even so, converting open-ended answers into explicit criteria yields effective, well-aligned reward signals.

Between the two rubric-guided methods, RaR-Implicit attains the strongest results overall; fixed weighted sums in RaR-Explicit offer more control but can be brittle. Explicit weighting can be difficult to tune but offers greater interpretability; we view the choice as application-dependent and leave it to practitioners. Future work could explore learned or dynamic weighting strategies that maintain interpretability while improving adaptability.

Rubrics enhance alignment with human preferences across model scales We evaluate alignment with humans by having LLM judges of varying sizes score chosen vs. rejected HealthBench-1k responses on a 1–10 scale under two settings: (i) rubric-guided (RaR-IMPLICIT), where the instance-specific rubric is provided, and (ii) rubric-free (DIRECT-LIKERT), where only the prompt and answers are shown. Figure 3 reports *pairwise preference accuracy* (the fraction of pairs where the preferred response receives the higher score). Rubric guidance improves accuracy for every judge size, with the largest gains for smaller judges, narrowing the gap to larger models. This indicates that explicit, context-specific criteria help judges distinguish subtle quality differences better than direct Likert scoring. Further analysis of judge-scale effects on GRPO training is detailed in Appendix A.9.

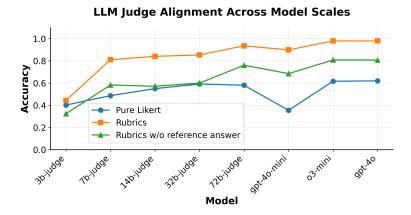


Figure 3: Alignment Study of LLM Judges across Model Scales. Rubrics as Rewards (orange) consistently improves alignment with human preferences across LLM judge sizes compared to direct Likert-based scoring (blue). Judge Alignment using synthetic rubrics without expert grounding (green) outperform the direct Likert baseline, but still fall short of expert-grounded rubrics (orange). The rubric structure especially benefits smaller judge models, helping them close the gap with larger models when guided by checklist-style criteria.

Training Method	Overall Score
Expert-Answer-SFT	20.4%
Simple-Likert	23.9%
Reference-Likert	31.7%
RaR-Implicit-Synthetic-NoRef	32.0%
RaR-Implicit-Synthetic	35.9%
RaR-Implicit-Human	34.8%

Table 1: Evaluation on HealthBench: Comparison of human- vs. synthetic-generated rubrics (with and without reference answers). RaR methods trained with GRPO significantly outperform Likertonly, Reference-based-likert and SFT baselines. Synthetic rubrics generated *without* access to reference answers perform notably worse, highlighting the importance of human-grounded guidance. Notably, human-authored rubrics and synthetic rubrics with access to references yield comparable performance.

Expert guidance is crucial for synthetic rubric generation Human guidance significantly influences the effectiveness of rubrics in capturing subtle human preferences. Figure 3 highlights performance differences between rubric-based evaluations that include reference answers and those without them. The data shows that rubrics developed with reference answers achieve higher accuracy, emphasizing that human insights integrated during rubric generation enable granular criteria and improved alignment with human preferences.

6 ABLATIONS

Impact of Rubric Generation Strategies in Real-World Domains How does the method of rubric generation affect downstream training in challenging, real-world settings? To investigate this, we hold out HealthBench-1k for evaluation and use 3.5k prompts from the remaining HealthBench pool to generate rubrics for training as it has access to human generated rubrics. Results are summarized in Table 1.

In-domain testing on HealthBench-1k amplifies RaR's gains: every instance-specific rubric-based method outperforms rubric-free baselines. Notably, even the weakest RaR variant significantly surpasses Reference-Likert, underscoring the advantage of structured supervision in subjective, open-ended domains like healthcare. We attribute this to the finer granularity and clarity rubrics pro-

3	/	ö
3	7	9
3	8	0
2	0	4

Ablation Setting	Overall Score
Essential-Only Rubrics	34.9%
No Categorical Labels	38.8%
No Pitfall Criteria	37.2%
All Rubrics	37.2%

Table 2: Ablation results for elements of rubric design on HealthBench-1k (trained on HealthBench-3.5k subset with Qwen2.5-7B base policy). Rubrics generated using o3-mini with access to reference answers.

Rubric Generation Model	Overall Score
O3-mini (Rubrics with reference)	35.9%
GPT-4o	34.2%
GPT-4o-mini	32.7%
O3-mini	32.4%
Qwen-72B-Instruct	32.7%
Qwen-32B-Instruct	31.1%
Qwen-7B-Instruct	31.9%

Table 3: Policy performance on HealthBench-1k when trained with GRPO using rubrics generated by different LLMs *without* reference answers. GPT-4o-generated rubrics yield the strongest performance, though they still fall short of rubrics generated with expert (reference-guided) supervision. Smaller aligned models (e.g., GPT-4o-mini, O3-mini) remain competitive with larger open-weight models, underscoring the importance of alignment and reasoning ability in rubric generation.

vide in assigning rewards-especially when correctness is not binary and answers may vary in tone, completeness, or safety relevance.

Moreover, we find that rubric quality is crucial: synthetic rubrics generated with reference-answer guidance consistently outperform those generated without it. This highlights the importance of incorporating expert signal, whether via human-in-the-loop annotations or high-quality reference completions, for generating effective and aligned rubrics. Purely synthetic rubrics, while scalable, currently fall short in capturing the subtle criteria required for robust training in high-stakes domains.

Elements of Rubric Design This ablation study examines how the structure and weighting of synthetic rubrics affect downstream performance on HealthBench-1k. As shown in Table 2, rubrics that include a broader range of criteria outperform those limited to essential checks, suggesting that richer evaluation signals lead to better learning. Interestingly, we observe minimal performance differences when including rubric weights or pitfall criteria during training. One possible explanation is that synthetically generating effective pitfall criteria is inherently difficult, as it requires anticipating the most common or critical failure modes of the model, a task that often demands human intuition and domain expertise. As a result, these synthetic negative criteria may lack the specificity or relevance needed to meaningfully penalize undesirable responses.

Impact of LLM Expertise on Rubric Quality To assess how the capabilities of rubric-generating LLMs affect downstream performance, we generate synthetic rubrics without access to reference answers and use them to train policies on HealthBench. This isolates the effect of LLM quality on reference-free rubric utility. Specifically, we evaluate on the HealthBench-1k subset, using models trained on rubrics generated for the remaining 4k training examples from HealthBench.

As shown in Table 3, larger or more capable LLMs generally produce more effective rubrics, with GPT-40 yielding the best performance among reference-free models. However, all remain below the performance of rubrics generated with reference guidance (e.g., O3-mini with access to reference answers). Additionally, model attributes such as instruction tuning and reasoning capabilities play a key role in the effectiveness of rubric generation.

7 RELATED WORK

RLVR across domains Reinforcement learning with verifiable rewards (RLVR) is expanding beyond math and code. GENERAL-REASONER trains on a 200k mixed corpus spanning physics, finance, and policy, and reports a ten-point gain on MMLU-Pro after GRPO fine-tuning (Ma et al., 2025). A follow-up extends RLVR to medicine, chemistry, psychology, and economics, showing that a single cross-domain reward model can supervise all four without task-specific tweaks (Su et al., 2025a). In healthcare, MED-RLVR applies similar methods to multiple-choice clinical QA, improving accuracy over supervised baselines while eliciting chain-of-thought from a 3B base model (Zhang et al., 2025). These results indicate steady progress, yet sparse signals, verifier reliability, and limited benchmark coverage remain open challenges.

Rubrics for evaluation and training Task-specific rubrics are increasingly used to evaluate LLMs in difficult-to-verify domains (Arora et al., 2025; Ruan et al., 2025; Hashemi et al., 2024; Pathak et al., 2025). Pathak et al. show that rubric-prompted LLM graders are more accurate and consistent than a question-agnostic checklist (Pathak et al., 2025). HEALTHBENCH scales this idea in medicine, pairing 48k clinician-written criteria with GPT-4 judges to score various axes (Arora et al., 2025). Beyond evaluation, rubrics are used to condition preference pairs for DPO (CPT; (Gallego, 2025)) and to guide checklist-based preference tuning in safety, instruction-following, and creative-writing settings (Viswanathan et al., 2025; Dineen et al., 2025; Kim et al., 2025). These lines of work primarily use rubrics to grade outputs or to condition preference data, often in non-reasoning domains such as safety, instruction following, or creative writing. In contrast, we use rubric criteria directly as reward signals for on-policy RL in expert-reasoning and real-world domains.

Learning from feedback signals RLHF trains policies with large numbers of human comparisons, which introduces subjectivity and can lead to reward hacking (Ouyang et al., 2022). RLVR reduces these issues by using programmatic checks, from exact matches on GSM8K and MATH to mixed-domain verifiers in GENERAL-REASONER and CROSS-DOMAIN RLVR (Ma et al., 2025; Su et al., 2025a), although signals can be sparse. Process supervision (Lightman et al., 2023) provides denser guidance via step-level labels, and MCTS-generated annotations or generative reward models such as THINKPRM improve performance, but with high annotation cost (Li et al., 2025; Khalifa et al., 2025). Rubric-based RL finds a middle ground by turning multiple rubric criteria into structured verifiers and using their scalar scores as denser rewards.

8 Conclusion

We introduced **Rubrics as Rewards** (**RaR**), a framework for post-training language models using structured, checklist-style rubrics as reward signals. By decomposing response evaluation into transparent, multi-criteria objectives—both subjective and objective—RaR provides a modular and interpretable alternative to preference-based methods. Our experiments demonstrate that rubric-guided training achieves strong performance across domains, significantly outperforming Likert-based baselines and matching or exceeding the performance of reference-based reward generation. Additionally, we show that RaR improves alignment with human preferences while reducing dependence on large judge models.

9 LIMITATIONS AND FUTURE WORK

Our work focuses on medicine and science to enable controlled experiments. This choice allows us to run controlled experiments, but broader validation across dialogue, tool use, or other agentic tasks remains an important direction. We evaluate only two reward aggregation strategies, implicit and explicit, since they capture complementary extremes of flexibility and control; future work could explore more advanced ways of combining rubric criteria, such as learning continuous weights for each criterion or dynamically adjusting weights over the course of training to mimic a curriculum (e.g., prioritizing essential correctness early, then gradually emphasizing more subtle qualities like style or safety). Finally, we use off-the-shelf LLMs as judges for accessibility and reproducibility; exploring specialized evaluators with stronger reasoning or generative reward models may yield further gains.

REPRODUCIBILITY STATEMENT

We document data construction and rubric synthesis in Sections 2 and 3, with full prompt templates in Appendices A.6, A.8, and A.10. The *RaR-Medicine* and *RaR-Science* training sets, including per-prompt rubrics, are provided as supplementary material. Training settings for on-policy GRPO (model, batch size, schedule, hardware) appear in Section 4.2 and Table 10. All judge prompts used to compute rewards for rubric-guided and Likert baselines are in Appendix A.7. Evaluation protocols, metrics, and prompts are detailed in Section 4.5 and Appendix A.5; we report means with 95% confidence intervals and use the open-source simple-evals framework. Trained models will be released upon publication.

REFERENCES

- David Anugraha, Zilu Tang, Lester James V Miranda, Hanyang Zhao, Mohammad Rifqi Farhansyah, Garry Kuwanto, Derry Wijaya, and Genta Indra Winata. R3: Robust rubric-agnostic reward models. *arXiv preprint arXiv:2505.13388*, 2025.
- Rahul K Arora, Jason Wei, Rebecca Soskin Hicks, Preston Bowman, Joaquin Quiñonero-Candela, Foivos Tsimpourlas, Michael Sharman, Meghan Shah, Andrea Vallone, Alex Beutel, et al. Healthbench: Evaluating large language models towards improved human health. *arXiv preprint arXiv:2505.08775*, 2025.
- Junying Chen, Zhenyang Cai, Ke Ji, Xidong Wang, Wanlong Liu, Rongsheng Wang, Jianye Hou, and Benyou Wang. Huatuogpt-o1, towards medical complex reasoning with llms. *arXiv* preprint *arXiv*:2412.18925, 2024a.
- Lichang Chen, Chen Zhu, Davit Soselia, Jiuhai Chen, Tianyi Zhou, Tom Goldstein, Heng Huang, Mohammad Shoeybi, and Bryan Catanzaro. Odin: Disentangled reward mitigates hacking in rlhf. arXiv preprint arXiv:2402.07319, 2024b.
- Xiusi Chen, Gaotang Li, Ziqi Wang, Bowen Jin, Cheng Qian, Yu Wang, Hongru Wang, Yu Zhang, Denghui Zhang, Tong Zhang, et al. Rm-r1: Reward modeling as reasoning. *arXiv preprint arXiv:2505.02387*, 2025.
- Ganqu Cui, Lifan Yuan, Zefan Wang, Hanbin Wang, Wendi Li, Bingxiang He, Yuchen Fan, Tianyu Yu, Qixin Xu, Weize Chen, et al. Process reinforcement through implicit rewards. *arXiv preprint arXiv:2502.01456*, 2025.
- Jacob Dineen, Aswin RRV, Qin Liu, Zhikun Xu, Xiao Ye, Ming Shen, Zhaonan Li, Shijie Lu, Chitta Baral, Muhao Chen, et al. Qa-lign: Aligning llms through constitutionally decomposed qa. *arXiv* preprint arXiv:2506.08123, 2025.
- Víctor Gallego. Configurable preference tuning with rubric-guided synthetic data. *arXiv* preprint arXiv:2506.11702, 2025.
- General Reasoning. General reasoning, 2025. URL https://gr.inc/.
- Arnav Gudibande, Eric Wallace, Charlie Snell, Xinyang Geng, Hao Liu, Pieter Abbeel, Sergey Levine, and Dawn Song. The false promise of imitating proprietary llms. *arXiv preprint arXiv:2305.15717*, 2023.
- Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, et al. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*, 2025a.
- Jiaxin Guo, Zewen Chi, Li Dong, Qingxiu Dong, Xun Wu, Shaohan Huang, and Furu Wei. Reward reasoning model. *arXiv preprint arXiv:2505.14674*, 2025b.
- Helia Hashemi, Jason Eisner, Corby Rosset, Benjamin Van Durme, and Chris Kedzie. Llm-rubric: A multidimensional, calibrated approach to automated evaluation of natural language texts. *arXiv* preprint arXiv:2501.00274, 2024.

- Aaron Hurst, Adam Lerer, Adam P Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, et al. Gpt-4o system card. arXiv preprint arXiv:2410.21276, 2024.
 - Aaron Jaech, Adam Kalai, Adam Lerer, Adam Richardson, Ahmed El-Kishky, Aiden Low, Alec Helyar, Aleksander Madry, Alex Beutel, Alex Carney, et al. Openai o1 system card. *arXiv* preprint arXiv:2412.16720, 2024.
 - Muhammad Khalifa, Rishabh Agarwal, Lajanugen Logeswaran, Jaekyeom Kim, Hao Peng, Moontae Lee, Honglak Lee, and Lu Wang. Process reward models that think. *arXiv preprint arXiv:2504.16828*, 2025.
 - Seungone Kim, Juyoung Suk, Shayne Longpre, Bill Yuchen Lin, Jamin Shin, Sean Welleck, Graham Neubig, Moontae Lee, Kyungjae Lee, and Minjoon Seo. Prometheus 2: An open source language model specialized in evaluating other language models. *arXiv* preprint arXiv:2405.01535, 2024.
 - Zae Myung Kim, Chanwoo Park, Vipul Raheja, Suin Kim, and Dongyeop Kang. Toward evaluative thinking: Meta policy optimization with evolving reward models. *arXiv* preprint arXiv:2504.20157, 2025.
 - Nathan Lambert, Jacob Morrison, Valentina Pyatkin, Shengyi Huang, Hamish Ivison, Faeze Brahman, Lester James V Miranda, Alisa Liu, Nouha Dziri, Shane Lyu, et al. T\" ulu 3: Pushing frontiers in open language model post-training. *arXiv preprint arXiv:2411.15124*, 2024.
 - Shuangtao Li, Shuaihao Dong, Kexin Luan, Xinhan Di, and Chaofan Ding. Enhancing reasoning through process supervision with monte carlo tree search. *arXiv preprint arXiv:2501.01478*, 2025.
 - Hunter Lightman, Vineet Kosaraju, Yuri Burda, Harrison Edwards, Bowen Baker, Teddy Lee, Jan Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. Let's verify step by step. In *The Twelfth International Conference on Learning Representations*, 2023.
 - Dakuan Lu, Xiaoyu Tan, Rui Xu, Tianchu Yao, Chao Qu, Wei Chu, Yinghui Xu, and Yuan Qi. Scp-116k: A high-quality problem-solution dataset and a generalized pipeline for automated extraction in the higher education science domain. *arXiv preprint arXiv:2501.15587*, 2025.
 - Xueguang Ma, Qian Liu, Dongfu Jiang, Ge Zhang, Zejun Ma, and Wenhu Chen. General-reasoner: Advancing llm reasoning across all domains. *arXiv preprint arXiv:2505.14652*, 2025.
 - OpenAI o3-mini. Openai o3-min, 2025. URL https://openai.com/index/openai-o3-mini/.
 - Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. Training language models to follow instructions with human feedback. Advances in neural information processing systems, 35: 27730–27744, 2022.
 - Aditya Pathak, Rachit Gandhi, Vaibhav Uttam, Yashwanth Nakka, Aaryan Raj Jindal, Pratyush Ghosh, Arnav Ramamoorthy, Shreyash Verma, Aditya Mittal, Aashna Ased, et al. Rubric is all you need: Enhancing llm-based code evaluation with question-specific rubrics. *arXiv preprint arXiv:2503.23989*, 2025.
 - Jie Ruan, Inderjeet Nair, Shuyang Cao, Amy Liu, Sheza Munir, Micah Pollens-Dempsey, Tiffany Chiang, Lucy Kates, Nicholas David, Sihan Chen, et al. Expertlongbench: Benchmarking language models on expert-level long-form generation tasks with structured checklists. *arXiv* preprint arXiv:2506.01241, 2025.
 - Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, YK Li, Y Wu, et al. Deepseekmath: Pushing the limits of mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300*, 2024.
 - Prasann Singhal, Tanya Goyal, Jiacheng Xu, and Greg Durrett. A long way to go: Investigating length correlations in rlhf. *arXiv preprint arXiv:2310.03716*, 2023.

- Ved Sirdeshmukh, Kaustubh Deshpande, Johannes Mols, Lifeng Jin, Ed-Yeremai Cardona, Dean Lee, Jeremy Kritz, Willow Primack, Summer Yue, and Chen Xing. Multichallenge: A realistic multi-turn conversation evaluation benchmark challenging to frontier llms. *arXiv preprint* arXiv:2501.17399, 2025.
 - Yi Su, Dian Yu, Linfeng Song, Juntao Li, Haitao Mi, Zhaopeng Tu, Min Zhang, and Dong Yu. Crossing the reward bridge: Expanding rl with verifiable rewards across diverse domains. *arXiv* preprint arXiv:2503.23829, 2025a.
 - Yi Su, Dian Yu, Linfeng Song, Juntao Li, Haitao Mi, Zhaopeng Tu, Min Zhang, and Dong Yu. Expanding rl with verifiable rewards across diverse domains. *arXiv preprint arXiv:2503.23829*, 2025b.
- Vijay Viswanathan, Yanchao Sun, Shuang Ma, Xiang Kong, Meng Cao, Graham Neubig, and Tongshuang Wu. Checklists are better than reward models for aligning language models. *arXiv* preprint *arXiv*:2507.18624, 2025.
- Haoxiang Wang, Yong Lin, Wei Xiong, Rui Yang, Shizhe Diao, Shuang Qiu, Han Zhao, and Tong Zhang. Arithmetic control of llms for diverse user preferences: Directional preference alignment with multi-objective rewards. *arXiv preprint arXiv:2402.18571*, 2024.
- Chenxi Whitehouse, Tianlu Wang, Ping Yu, Xian Li, Jason Weston, Ilia Kulikov, and Swarnadeep Saha. J1: Incentivizing thinking in llm-as-a-judge via reinforcement learning. *arXiv* preprint arXiv:2505.10320, 2025.
- Zihuiwen Ye, Fraser Greenlee-Scott, Max Bartolo, Phil Blunsom, Jon Ander Campos, and Matthias Gallé. Improving reward models with synthetic critiques. *arXiv preprint arXiv:2405.20850*, 2024.
- Weizhe Yuan, Jane Yu, Song Jiang, Karthik Padthe, Yang Li, Ilia Kulikov, Kyunghyun Cho, Dong Wang, Yuandong Tian, Jason E Weston, et al. Naturalreasoning: Reasoning in the wild with 2.8 m challenging questions. *arXiv preprint arXiv:2502.13124*, 2025.
- Sheng Zhang, Qianchu Liu, Guanghui Qin, Tristan Naumann, and Hoifung Poon. Med-rlvr: Emerging medical reasoning from a 3b base model via reinforcement learning. *arXiv* preprint arXiv:2502.19655, 2025.
- Xuandong Zhao, Zhewei Kang, Aosong Feng, Sergey Levine, and Dawn Song. Learning to reason without external rewards. *arXiv* preprint arXiv:2505.19590, 2025.
- Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, et al. Judging llm-as-a-judge with mt-bench and chatbot arena. *Advances in neural information processing systems*, 36:46595–46623, 2023.

A APPENDIX

A.1 USE OF LLMS

We use LLMs for minor editorial support, including grammar, spelling, punctuation, latex formatting and word choice corrections, with usage limited to enhancing readability and clarity.

A.2 DETAILS OF RAR-MEDICINE DATASET

The following is illustrates an example from the RaR-Medicine dataset with more details about statistics and categorical distribution of rubrics listed in Table 4 and 6.

RaR-Medicine Example

Question. A 50-year-old male patient weighs 65 kg with a pH of 7.05, PCO2 of 15 mmHg, HCO3 of 5 mEq/L, and a base deficit of -40 mEq/L. How much sodium bicarbonate should be administered in the first 4 hours to correct his metabolic acidosis?

Reference Answer. To address the severe metabolic acidosis in this 50-year-old patient (weight 65 kg) with a pH of 7.05, it is crucial to approach treatment with caution to avoid complications from overcorrection. The base deficit is -40 mEq/L, which indicates a significant acid-base disturbance. Typically, the formula for bicarbonate administration is:

Bicarbonate required = Base Deficit(mEq/L) \times Body Weight(kg) \times 0.3

Applying the values:

$$40 \,\mathrm{mEq/L} \times 65 \,\mathrm{kg} \times 0.3 = 780 \,\mathrm{mEq}$$

However, it is standard practice to administer a partial correction initially, aiming for safety and gradual improvement. Therefore, administering about 150 mEq of sodium bicarbonate over the first 4 hours is a prudent approach. This amount aligns with safe clinical guidelines for initial treatment in severe metabolic acidosis, ensuring the patient starts on the path to stabilization without risking harm from rapid correction.

Rubric (checklist and weights).

- **Bicarbonate Calculation** (*w*=5): Essential Criteria: The response must correctly identify and apply the formula (Base Deficit × Body Weight × 0.3) to determine the bicarbonate requirement.
- Safe Dosing Recommendation (w=5): Essential Criteria: The response must state a clear recommendation of administering about 150 mEq of sodium bicarbonate over the first 4 hours.
- Partial Correction Justification (w=4): Important Criteria: The response should explain that only a partial correction is administered initially to avoid complications from rapid overcorrection.
- Step-by-Step Calculation (w=3): Important Criteria: The response must detail the calculation steps, showing that 40 mEq/L × 65 kg × 0.3 equals 780 mEq, before noting the adjusted dose for safety.
- Base Deficit Interpretation (w=2): Optional Criteria: The response may mention that the base deficit of -40 mEq/L indicates severe metabolic acidosis requiring cautious treatment.
- Patient Data Accuracy (w=3): Important Criteria: The response must accurately incorporate the patient's weight of 65 kg and his critical pH, PCO2, and HCO3 values into the explanation.
- Avoid Overcorrection Risk (w=-1): Pitfall Criteria: Does not mention the risks associated with rapid overcorrection of metabolic acidosis if only the full calculated bicarbonate amount is administered.

Metric	Value
Total examples	20,166
Avg. rubrics per question	7.5
Avg. question length (words)	45.0

Table 4: Aggregate statistics for the RaR-Medicine dataset (train and validation) dataset.

7	702	
	703	
	704	
	705	
	706	
	707	
	708	
	709	
	710	
	711	
	712	
	713	
	714	
7		
7		
	717	
	718 710	
7		
7		
7		
7		
7		
7		
7		
7		
7		
7		
7		
7		
7		
7		
7		
7		
7		
7	737	
7	738	
7		
7		
7		
7		
7		
7		
7		
7		
7		
7	748	
7	749	
7	750	
7	751	
7	752	
-	700	

754

755

Rubric Type	Count	Percent
Important	52,748	34.1
Essential Optional	47,584 34,261	30.7 22.1
Pitfall	20,215	13.1

Table 5: Rubric-type distribution across all 20,166 examples.

Topics	Count	Percent
Total examples	20,166	100.0
Medical Diagnosis	10,147	50.3
Medical Treatment	3,235	16.0
Medical Knowledge	2,557	12.7
Medical Diag. and Mngmnt	2,033	10.1
Medical Biology	770	3.8
Other	428	2.1
Medical Ethics	377	1.9
Health Physics	276	1.4
Epidemiology & Pub. Health	216	1.1
General Medicine	113	0.6
Forensic Medicine	14	0.1

Table 6: Distribution of topics in the medical training and validation dataset

A.3 DETAILS OF RAR-SCIENCE DATASET

This section shows an illustrative example from the RaR-Science dataset with more details about statistics and categorical distibution of rubrics listed in Table 7, 8, 9.

RaR-Science Example

Question. Determine the solubility of boric acid (H_3BO_3) in ethanol (C_2H_5OH) compared to its solubility in benzene (C_6H_6) , considering the principles of 'likes dissolve likes' and the role of K_{sp} values. Explain your reasoning and provide examples of how immulcifiers could affect the solubility of substances in different solvents.

Reference Answer. Boric acid is more soluble in ethanol than in benzene.

Rubric (checklist and weights).

- **Correct Solubility Direction** (*w*=5): Essential Criteria: The response must clearly identify that boric acid is more soluble in ethanol than in benzene.
- **Polarity Principle** (*w*=5): Essential Criteria: The answer should explain how the 'like dissolves like' principle applies by contrasting the polar nature of ethanol with the non-polar character of benzene.
- **Ksp Context** (*w*=4): Important Criteria: The response should account for the role of Ksp values, discussing their typical relevance to solubility even though boric acid is a covalent compound rather than an ionic one.
- Immulcifier Explanation (*w*=4): Important Criteria: The answer should explain how immulcifiers could modify solubility, providing an example of their effect on solvation in different solvents.
- Chemical Properties (w=4): Important Criteria: The response should analyze the inherent chemical properties of boric acid and the solvents to justify the observed solubility differences.
- Avoid Ionic Assumptions (w=-1): Pitfall Criteria: The answer must not incorrectly assume that Ksp values for ionic compounds directly determine the solubility of a covalent acid such as boric acid.
- Enhanced Detail (*w*=2): Optional Criteria: The response may include additional examples or a concise explanation of solvation dynamics to further illustrate how solubility is influenced.

Metric	Value
Total examples	20,625
Avg. rubrics per question	7.5
Avg. question length (words)	52.6

Table 7: Aggregate statistics for the full medical (train and validation) dataset.

Rubric Type	Count	Percent
Important Essential Optional Pitfall	52,315 42,739 33,622 21,808	34.8 28.4 22.3 14.5

Table 8: Rubric-type distribution across all 20 625 examples.

A.4 TRAINING DETAILS

The training hyperparameters are described in Table 10.

Topics	Count	Percent
Total examples	20625	100.0
General Chemistry	3163	15.3
Quantum Mechanics	3158	15.3
Physical Chemistry	2761	13.4
Statistical Mechanics	2530	12.3
Organic Chemistry	2059	10.0
General Physics	1439	7.0
Condensed Matter Physics	1387	6.7
Genetics	1378	6.7
Molecular Biology	815	4.0
Astrophysics	409	2.0
Inorganic Chemistry	407	2.0
Analytical Chemistry	398	1.9
Electromagnetism	239	1.2
Optics	143	0.7
High Energy Physics	116	0.6
Electromagnetic Theory	105	0.5
Electromagnetics	72	0.3
Relativistic Mechanics	46	0.2

Table 9: Distribution of topics in the STEM training and validation dataset

Rubric-based evaluation consistently yields stronger policies across all judge sizes. The most pronounced improvement appears with Qwen-7B-Instruct (+0.047), where rubric guidance lifts it from weakest to nearly matching larger models. Additionally, rubric-based scores are more tightly clustered (0.250–0.279) than those from Likert-only judges (0.220–0.254), indicating improved consistency.

These results suggest that rubrics help smaller judges approximate high-quality supervision by breaking evaluation into interpretable, binary criteria. This structured approach mitigates scale-related limitations, enabling more reliable reward modeling even with limited-capacity evaluators.

0	4	n
8	1	1
	1	
	1	
	1	
	1	
	1	
8		7
	1	
	1	
	2	
	2	
	2	
8	2	3
8	2	4
8	2	5
8	2	6
8	2	7
8	2	8
8	2	9
8	3	0
8	3	1
8	3	2
	3	
	3	
	3	
	3	
	3	
	3	
	3	
	о 4	
	4	
	4	
	4	
	4	
	4	
	4	
	4	
8	4	8
8	4	9
8	5	0
8	5	1
8	5	2
	5	
	5	
	5	
	5	
	5 5	
	ວ 5	
ď	5	y

862

Hyperparameters		
num_rollouts_per_prompt	16	
batch_size (effective)	96	
sampling_temperature	1.0	
warmup_ratio	0.1	
learning_rate	5.0e-06	
lr_scheduler_type	constant_with_warmup	
max_length	3584	
num_train_steps	300	

Table 10: GRPO hyperparameter settings for Medical and Science domains.

Judge Model	RaR-Implicit	Direct-Likert
GPT-4o-mini	27.9%	25.3%
Qwen-32B-Instruct	26.2%	25.4%
Qwen-14B-Instruct	25.0%	24.9%
Qwen-7B-Instruct	26.7%	22.0%

Table 11: Judge quality comparison: rubric-based evaluation vs pure Likert scoring on synthetic medical rubrics.

A.5 EVALUATION PROMPTS

GPQA Evaluation Prompt

Determine whether the following model response matches the ground truth answer.

```
## Ground truth answer##: Option {correct_answer} or {correct_answer_text}
## Model Response ##: {response_text}
```

A response is considered correct if it's final answer is the correct option letter (A, B, C, or D), or has the correct answer text.

Please respond with only "Yes" or "No" (without quotes). Do not include a rationale.

A.6 PREDEFINED STATIC RUBRICS

Predefined Static Rubrics for RaR-Static Method

- The response contains correct information without factual errors, inaccuracies, or hallucinations that could mislead the user.
- The response fully answers all essential parts of the question and provides sufficient detail where needed.
- The response is concise and to the point, avoiding unnecessary verbosity or repetition.
- The response effectively meets the user's practical needs, provides actionable information, and is genuinely helpful for their situation.

A.7 LLM-JUDGE PROMPTS

Prompt for RAR-IMPLICIT Method

System Prompt:

You are an expert evaluator. Given a user prompt, a generated response, and a list of quality rubrics, please rate the overall quality of the response on a scale of 1 to 10 based on how well it satisfies the rubrics. Consider all rubrics holistically when determining your score. A response that violates multiple rubrics should receive a lower score, while a response that satisfies all rubrics should receive a higher score. Start your response with a valid JSON object that starts with """ json" and ends with """." The JSON object should contain a single key "rating" and the value should be an integer between 1 and 10.

Example response:

```
```json
{
"rating": 7
}```
```

#### **User Prompt Template:**

Given the following prompt, response, and rubrics, please rate the overall quality of the response on a scale of 1 to 10 based on how well it satisfies the rubrics.

```
<prompt>
{prompt}
</prompt>
</prompt>
<response>
{response}
</response>
<rubrics>
{rubrics>
{rubrics>
Your JSON Evaluation:
```

## 

#### **Prompt for DIRECT-LIKERT Baseline**

#### **System Prompt:**

You are an expert evaluator. Given a user prompt and a generated response, please rate the overall quality of the response on a scale of 1 to 10, where 1 is very poor and 10 is excellent.

Start your response with a valid JSON object that starts with "``json" and ends with "``". The JSON object should contain a single key "rating" and the value should be an integer between 1 and 10.

Example response:

```
```json
{
"rating": 8
```

User Prompt Template:

Given the following prompt, and response, please rate the overall quality of the response on a scale of 1 to 10.

```
<prompt>
{prompt}
</prompt>
</response>
{response}
</response>
```

Your JSON Evaluation:

Prompt for REFERENCE-LIKERT Baseline

System Prompt:

You are an expert evaluator. Given a user prompt, a reference response, and a generated response, please rate the overall quality of the generated response on a scale of 1 to 10 based on how well it compares to the reference response.

Consider factors such as accuracy, completeness, coherence, and helpfulness when comparing to the reference. The reference response represents a high-quality answer that you should use as a benchmark. Start your response with a valid JSON object that starts with """ json" and ends with """.". The JSON object should contain a single key "rating" and the value should be an integer between 1 and 10. Example response: ""json"

```
{
"rating": 8
}```
```

User Prompt Template: Given the following prompt, reference response, and generated response, please rate the overall quality of the generated response on a scale of 1 to 10 based on how well it compares to the reference.

```
<prompt>
{prompt}
</prompt>
</prompt>
<reference_response>
{reference}
</reference_response>
<generated_response>
{response}
</generated_response>
```

Your JSON Evaluation:

A.8 SYNTHETIC PREFERENCE SET GENERATION

We leverage the publicly-released HEALTHBENCH (Arora et al., 2025) corpus, which contains **5,000** health-related prompts accompanied by expert-written answers. Of these, **4,203** datapoints already include an *ideal* completion vetted by licensed clinicians. For every such prompt—ideal pair we automatically generate a *perturbed* counterpart using o3 with the structured template shown below. The template forces the model to (i) spell out a <code>[reasoning]</code> plan for degrading quality, (ii) emit the degraded <code>[perturbed_completion]</code>, and (iii) log exact <code>[chunks_added]</code> and <code>[chunks_removed]</code>. Perturbations are accepted only after manual screening confirms that they are *objectively worse*, along at least one axis of medical accuracy, completeness, clarity, safety, specificity, structure, or tone, while remaining coherent and free of dangerous advice. We further exclude the prompts from HealthBench-1k used for ablations. This procedure produces a balanced evaluation set of **3,027 preferred** and **3,027 perturbed** responses (6,054 total), which we use in the rubric-versus-Likert experiments in Section 5. The prompt used for this generation is detailed in Figure A.10.

A.9 JUDGE QUALITY IMPACTS ON POST-TRAINING

We assess whether rubric-guided evaluation improves judge effectiveness compared to rubric-free Likert scoring when used for GRPO training. Table 11 reports judge accuracy on synthetic medical data, with all policies trained using Qwen2.5-7B and varying judge models.

975

976

977

978

979

980

981

982 983

984

985

986

987

988

990

991

992

993

994

995

996

997 998

999

1000

1001

1002

1003

1004

1008

1009

1010 1011

1012

1013

1014

1015

1016

1017

1020

Prompt for Synthetic Rubrics Generation: Medical Domain

You are an expert rubric writer. Your job is to generate a self-contained set of evaluation criteria ("rubrics") for judging how good a response is to a given question. Rubrics can cover aspects of a response such as, but not limited to, factual correctness, ideal-response characteristics, style, completeness, helpfulness, harmlessness, patient-centeredness, depth of reasoning, contextual relevance, and empathy. Each item must be self-contained – non expert readers should not need to infer anything or consult external information. Begin each description with its category: "Essential Criteria: ...", "Important Criteria: ...", "Optional Criteria: ...", or "Pitfall Criteria: Does not mention ...". Inputs:

- question: The full question text.
- reference_answer: The ideal answer, including any specific facts, explanations, or advice.

Total items:

• Choose 7–20 rubric items based on the complexity of the question.

Each rubric item:

- **title** (2–4 words).
- **description**: One sentence starting with its category prefix that explicitly states exactly what to look for. For example:
 - Essential Criteria: Identifies non-contrast helical CT scan as the most sensitive modality for ureteric stones.
 - Pitfall Criteria: Does not mention identifying (B) as the correct answer.
 - Important Criteria: Explains that non-contrast helical CT detects stones of varying sizes and compositions.
 - Optional Criteria: States "The final answer is (B)" or similar answer choice formatting.
- weight: For Essential/Important/Optional, use 1-5 (5 = most important); for Pitfall, use -1 or -2.

Category guidance:

- Essential: Critical facts or safety checks; if missing, the response is invalid (weight 5).
- Important: Key reasoning, completeness, or clarity; strongly affects quality (weight 3–4).
- **Optional**: Helpful style or extra depth; nice to have but not deal-breaking (weight 1–2).
- **Pitfall**: Common mistakes or omissions specific to this prompt—identify things a respondent often forgets or misstates. Each Pitfall description must begin with "Pitfall Criteria: Does not mention ..." or "Pitfall Criteria: Recommends ..." and use weight -1 or -2.

To ensure self-contained guidance:

- When referring to answer choices, explicitly say "Identifies (A)", "Identifies (B)", etc., rather than vague phrasing.
- If the format requires a conclusion like "The final answer is (B)", include a rubric item such as:
 - Essential Criteria: Includes a clear statement "The final answer is (B)".
- If reasoning should precede the answer, include a rubric like:
 - Important Criteria: Presents the explanation before stating the final answer.
- If brevity is valued, include a rubric like:
 - Optional Criteria: Remains concise and avoids unnecessary detail.
- If the question context demands mention of specific findings, include that explicitly (e.g., "Essential Criteria: Mentions that CT does not require contrast").

Output: Provide a JSON array of rubric objects. Each object must contain exactly three keys—title, description, and weight. Do not copy large blocks of the question or reference_answer into the text. Each description must begin with its category prefix, and no extra keys are allowed.

Now, given the question and reference answer, generate the rubric as described. The reference answer is an ideal response but not necessarily exhaustive; use it only as guidance.

1067

1074

1078

Prompt for Synthetic Rubric Generation: Science Domain

You are an expert rubric writer for science questions in the domains of Biology, Physics, and Chemistry. Your job is to generate a self-contained set of evaluation criteria ("rubrics") for judging how good a response is to a given question in one of these domains. Rubrics can cover aspects such as factual correctness, depth of reasoning, clarity, completeness, style, helpfulness, and common pitfalls. Each rubric item must be fully self-contained so that non-expert readers need not consult any external information.

Inputs:

- question: The full question text.
- reference_answer: The ideal answer, including any key facts or explanations.

Total items:

• Choose 7–20 rubric items based on question complexity.

Each rubric item must include exactly three keys:

- 1. **title** (2–4 words)
- 2. **description**: One sentence beginning with its category prefix, explicitly stating what to look for. For example:
 - Essential Criteria: States that in the described closed system, the total mechanical energy (kinetic plus potential) before the event equals the total mechanical energy after the event.
 - Important Criteria: Breaks down numerical energy values for each stage, demonstrating that initial kinetic energy plus initial potential energy equals final kinetic energy plus final potential energy.
 - Optional Criteria: Provides a concrete example, such as a pendulum converting between kinetic and potential energy, to illustrate how energy shifts within the system.
 - Pitfall Criteria: Does not mention that frictional or air-resistance losses are assumed negligible when applying conservation of mechanical energy.
- 3. weight: For Essential/Important/Optional, use 1–5 (5 = most important); for Pitfall, use –1 or –2.

Category guidance:

- Essential: Critical facts or safety checks; omission invalidates the response.
- Important: Key reasoning or completeness; strongly affects quality.
- · Optional: Nice-to-have style or extra depth.
- Pitfall: Common mistakes or omissions; highlight things often missed.

Format notes:

- When referring to answer choices, explicitly say "Identifies (A)", "Identifies (B)", etc.
- If a clear conclusion is required (e.g. "The final answer is (B)"), include an Essential Criteria for it.
- If reasoning should precede the final answer, include an Important Criteria to that effect.
- If brevity is valued, include an Optional Criteria about conciseness.

Output: Provide a JSON array of rubric objects. Each object must contain exactly three keys—title, description, and weight. Do not copy large blocks of the question or reference_answer into the text. Each description must begin with its category prefix, and no extra keys are allowed.

Now, given the question and reference_answer, generate the rubric as described. The reference answer is an ideal response but not necessarily exhaustive; use it only as guidance.

A.10 PERTURBED DATASET GENERATION

Prompt for Health Response Perturbation

You are tasked with creating perturbed versions of health-related responses to generate training data for evaluating model robustness. Given an original health prompt and an ideal completion, you must generate a perturbed version that is strictly worse than the ideal response but still plausible.

Input Format:

 [prompt]: prompt

[ideal_completion]: ideal_completion

Task Instructions: For the perturbed version, follow this exact process:

[reasoning]: Explain your strategy for making the response worse. Consider these perturbation approaches:

- Medical Accuracy: Introduce subtle medical inaccuracies or outdated information - Completeness: Remove crucial information or context - Clarity: Make explanations less clear or more confusing - Safety: Reduce appropriate medical disclaimers or caveats - Specificity: Replace specific advice with vague generalizations - Structure: Disorganize the logical flow of information - Tone: Make the response less empathetic or more dismissive

[perturbed_completion]: Create the modified response based on your reasoning strategy.

[chunks_added]: List any new text segments you added (even if they're incorrect or unhelpful)

[chunks_removed]: List any text segments you removed from the ideal response

Quality Guidelines:

- Ensure Degradation: Each perturbed response must be objectively worse than the ideal response - Maintain Plausibility: The perturbed responses should still seem like reasonable attempts (not obviously nonsensical) - Vary Approaches: Use different perturbation strategies for each version - Track Changes Precisely: Be specific about what text was added or removed - Health Context Sensitivity: Be mindful that health misinformation can be harmful - focus on subtle degradations rather than dangerous advice

Example Perturbation Strategies:

- Remove important safety warnings or medical disclaimers - Add outdated medical information that was once accepted but is now known to be incorrect - Remove specific dosage information or timing details - Add overly general statements that replace specific guidance - Remove context about when to seek professional medical help - Add confusing or contradictory information - Remove step-by-step instructions and replace with vague advice - Add unnecessarily complex medical jargon without explanation

A.11 MODEL AND BASELINE RESPONSE EXAMPLES

Example from RaR-Implicit

System Prompt: You are a helpful assistant.

User Prompt: In a parallel universe where a magnet can have an isolated North or South pole, Maxwell's equations look different. But, specifically, which of those equations are different?

A) The one related to the divergence of the magnetic field. B) The ones related to the circulation of the electric field and the divergence of the magnetic field. C) The one related to the circulation of the magnetic field and the flux of the electric field. D) The ones related to the divergence and the curl of the magnetic field. Please reason step by step, and provide your final answer choice at the end of your response in the format "The final answer is: LETTER" where LETTER is one of A/B/C/D.

Assistant:

To approach this question, let's first recall Maxwell's equations in the standard form that apply to our familiar universe where magnetic poles always appear in pairs (North and South):

1. **Gauss's Law for Electricity:**

$$\nabla \cdot \mathbf{E} = \frac{\rho}{\epsilon_0}$$

This equation relates the electric field divergence to the charge density.

2. **Gauss's Law for Magnetism:**

$$\nabla \cdot \mathbf{B} = 0$$

This equation states that there are no isolated magnetic charges (magnetic monopoles).

3. **Faraday's Law of Induction:**

$$\nabla \times \mathbf{E} = -\frac{\partial \mathbf{B}}{\partial t}$$

This equation describes how a time-varying magnetic field induces an electric field.

4. **Ampère's Law with Maxwell's Addition:**

 $\nabla \times \mathbf{B} = \mu_0 \mathbf{J} + \mu_0 \epsilon_0 \frac{\partial \mathbf{E}}{\partial t}$

This equation relates the curl of the magnetic field to the current density and the time-varying electric field. Now, let's consider a universe where isolated magnetic poles (monopoles) exist. In such a scenario, the fundamental property that changes is the existence of magnetic charge, which means Gauss's Law for Magnetism must be modified to account for these monopoles. Let's denote the magnetic charge density by m.

The modified Gauss's Law for Magnetism in the presence of monopoles becomes:

 $abla \cdot \mathbf{B} = rac{\mathbf{m}}{\mu_0}$

This equation now includes the effect of magnetic charges, allowing for a non-zero divergence of the magnetic field.

Let's analyze the options given:

- A) **The one related to the divergence of the magnetic field.** In the presence of magnetic monopoles, the divergence of the magnetic field is no longer zero. This directly affects Gauss's Law for Magnetism, which is one of Maxwell's equations. Therefore, this equation is different.
- B) **The ones related to the circulation of the electric field and the divergence of the magnetic field.** Faraday's Law of Induction $(\nabla \times \mathbf{E} = -\frac{\partial \mathbf{B}}{\partial t})$ and Gauss's Law for Magnetism $(\nabla \cdot \mathbf{B} = 0)$ are the ones that involve the circulation of the electric field and the divergence of the magnetic field, respectively. Since Gauss's Law for Magnetism changes, this option is correct.
- C) **The one related to the circulation of the magnetic field and the flux of the electric field.** This option refers to Ampère's Law with Maxwell's addition, which involves the curl of the magnetic field and the time-varying electric field. The presence of magnetic monopoles does not directly affect this equation, as it primarily deals with the dynamics between the electric and magnetic fields. Therefore, this equation remains unchanged.
- D) **The ones related to the divergence and the curl of the magnetic field.** This option encompasses both Gauss's Law for Magnetism and Ampère's Law with Maxwell's addition. We have already established that Gauss's Law for Magnetism is different, but Ampère's Law remains the same in terms of its fundamental form. Thus, not all equations in this category are different.
- Given this analysis, the equations that are different in a universe with isolated magnetic poles are those related to the divergence of the magnetic field and the circulation of the electric field.

Therefore, the final answer is: |B|.

Results:

Correct Answer: B Extracted Answer: B

1167 Score: 1.0

Example from RaR-Implicit

System Prompt: You are a helpful assistant.

User Prompt: Two quantum states with energies E1 and E2 have a lifetime of 10^{-9} sec and 10^{-8} sec, respectively. We want to clearly distinguish these two energy levels. Which one of the following options could be their energy difference so that they can be clearly resolved?

A) 10^{-9} eV

B) 10^{-11} eV

C) 10^{-4} eV

D) 10^{-8} eV

Please reason step by step, and provide your final answer choice at the end of your response in the format "The final answer is: LETTER" where LETTER is one of A/B/C/D.

Assistant:

To determine which energy difference allows us to clearly distinguish the two quantum states with lifetimes of 10^{-9} seconds and 10^{-8} seconds, we need to consider the principle of resolving two energy levels in quantum mechanics. The key concept here is the **uncertainty principle** and how it applies to the energy-time uncertainty relation. Let's break this down step by step:

Step 1: Understanding the Energy-Time Uncertainty Relation The energy-time uncertainty principle states that the product of the uncertainties in energy (ΔE) and time (Δt) must be greater than or equal to $\hbar/2$, where \hbar is the reduced Planck's constant ($\hbar = h/2\pi \approx 1.054 \times 10^{-34} \, \mathrm{J \, s}$):

$$\Delta E \cdot \Delta t \ge \frac{\hbar}{2}$$

Step 2: Relating Lifetime to Energy Uncertainty The lifetime (τ) of a quantum state is inversely related to the uncertainty in energy (ΔE) . For a state with lifetime τ , the uncertainty in energy is given by:

$$\Delta E = \frac{\hbar}{\tau}$$

Step 3: Applying the Relation to Given Lifetimes Let's apply this relation to the two given lifetimes: 1. For the state with $\tau_1 = 10^{-9}$ seconds:

$$\Delta E_1 = \frac{\hbar}{\tau_1} = \frac{1.054 \times 10^{-34} \,\mathrm{J \, s}}{10^{-9} \,\mathrm{s}} = 1.054 \times 10^{-25} \,\mathrm{J}$$

2. For the state with $\tau_2 = 10^{-8}$ seconds:

$$\Delta E_2 = \frac{\hbar}{\tau_2} = \frac{1.054 \times 10^{-34} \,\mathrm{J s}}{10^{-8} \,\mathrm{s}} = 1.054 \times 10^{-26} \,\mathrm{J}$$

Step 4: Converting Energy Uncertainty to Electronvolts To compare these values with the given options, we need to convert joules to electronvolts. Recall that $1\,\mathrm{eV} = 1.602 \times 10^{-19}\,\mathrm{J}$:

1. For ΔE_1 :

$$\Delta E_1 = 1.054 \times 10^{-25} \,\mathrm{J} \times \frac{1 \,\mathrm{eV}}{1.602 \times 10^{-19} \,\mathrm{J}} \approx 6.58 \times 10^{-7} \,\mathrm{eV}$$

2. For ΔE_2 :

$$\Delta E_2 = 1.054 \times 10^{-26} \,\mathrm{J} \times \frac{1 \,\mathrm{eV}}{1.602 \times 10^{-19} \,\mathrm{J}} \approx 6.58 \times 10^{-8} \,\mathrm{eV}$$

Step 5: Determining the Resolvable Energy Difference To clearly distinguish between the two energy levels, the energy difference (ΔE) must be at least as large as the larger of the two ΔE values calculated above. This ensures that the energy difference is larger than the uncertainty in energy for both states. From our calculations: - $\Delta E_1 \approx 6.58 \times 10^{-7} \, \mathrm{eV}$ - $\Delta E_2 \approx 6.58 \times 10^{-8} \, \mathrm{eV}$

The larger uncertainty is ΔE_1 , so we need:

$$\Delta E > 6.58 \times 10^{-7} \, \text{eV}$$

Step 6: Comparing with Given Options Now, let's compare this result with the provided options: A) 10^{-9} eV B) 10^{-11} eV C) 10^{-4} eV D) 10^{-8} eV

Clearly, 6.58×10^{-7} eV is closest to and larger than 10^{-8} eV.

Final Answer Thus, the energy difference that allows us to clearly distinguish the two energy levels is:

D

Results:

Correct Answer: C Extracted Answer: D

Score: 0.0

Example from Reference-Likert

System Prompt: You are a helpful assistant.

User Prompt: In a parallel universe where a magnet can have an isolated North or South pole, Maxwell's equations look different. But, specifically, which of those equations are different?

A) The one related to the divergence of the magnetic field. B) The ones related to the circulation of the electric field and the divergence of the magnetic field. C) The one related to the circulation of the magnetic field and the flux of the electric field. D) The ones related to the divergence and the curl of the magnetic field. Please reason step by step, and provide your final answer choice at the end of your response in the format "The final answer is: LETTER" where LETTER is one of A/B/C/D.

Assistant:

To address the question of how Maxwell's equations transform in a hypothetical universe where magnets can possess isolated North or South poles (contrary to our reality where magnetic monopoles are theorized but not observed), we need to carefully examine which of Maxwell's equations are fundamentally altered by the absence or presence of magnetic monopoles. Let's reason through this step-by-step:

1. **Understanding the Role of Magnetic Monopoles in Maxwell's Equations**

- **Traditional Maxwell's Equations (our universe):** These equations elegantly describe electromag-

- **Traditional Maxwell's Equations (our universe):** These equations elegantly describe electromagnetism, including the behavior of electric and magnetic fields, and they inherently assume the absence of magnetic monopoles (i.e., isolated North or South magnetic poles). The key equations here are: Gauss's Law for Magnetism: $\nabla \cdot \mathbf{B} = 0$, stating that magnetic fields have no sources (no magnetic monopoles). Ampère's Law with Maxwell's addition: $\nabla \times \mathbf{B} = \mu_0 \mathbf{J} + \mu_0 \epsilon_0 \frac{\partial \mathbf{E}}{\partial t}$, linking magnetic fields to electric currents and time-varying electric fields.
- -**Hypothetical Universe with Magnetic Monopoles:** If magnetic monopoles exist, Maxwell's equations must be adjusted to account for these new magnetic charges. The presence of magnetic monopoles implies: A modified Gauss's Law for Magnetism: $\nabla \cdot \mathbf{B} = 4\pi \mathbf{m}$, where \mathbf{m} represents the magnetic monopole density (analogous to electric charge density ρ in Gauss's Law for Electricity). Adjustments to Ampère's Law to include magnetic monopole currents: $\nabla \times \mathbf{B} = \mu_0(\mathbf{J} + \mathbf{m}) + \mu_0 \epsilon_0 \frac{\partial \mathbf{E}}{\partial t}$, explicitly incorporating both electric currents \mathbf{J} and magnetic monopole currents \mathbf{m} . ### 2. **Examining Each Maxwell Equation in Detail**

Given the above, let's examine which Maxwell equations are **different** (or altered) in a universe with magnetic monopoles:

- ** $\bar{\bf A}$) The one related to the divergence of the magnetic field ($\nabla \cdot {\bf B}$):** **Change:** In the presence of magnetic monopoles, Gauss's Law for Magnetism changes from $\nabla \cdot {\bf B} = 0$ to $\nabla \cdot {\bf B} = 4\pi {\bf m}$. This directly modifies how the divergence of the magnetic field ${\bf B}$ is related to magnetic monopole densities ${\bf m}$. **Conclusion:** The divergence of the magnetic field equation is **fundamentally different** due to the explicit inclusion of magnetic monopoles.
- -**B) The ones related to the circulation of the electric field and the divergence of the magnetic field:** **Change:** Ampère's Law with Maxwell's addition remains largely unchanged in form but its interpretation shifts: $\nabla \times \mathbf{B} = \mu_0 \mathbf{J} + \mu_0 \epsilon_0 \frac{\partial \mathbf{E}}{\partial t}$ now includes both electric currents \mathbf{J} and magnetic monopole currents \mathbf{m} . However, the core structure of Ampère's Law (involving $\nabla \times \mathbf{B}$) remains: **Electric field circulation:** $\nabla \times \mathbf{E} = -\frac{\partial \mathbf{B}}{\partial t}$ (Faraday's Law) is not directly altered by magnetic monopoles. **Magnetic field divergence:** As discussed, $\nabla \cdot \mathbf{B}$ changes significantly with magnetic monopoles. **Conclusion:** While Ampère's Law's form is not altered, its physical interpretation and inclusion of magnetic monopole currents (\mathbf{m}) make it **different** in a universe with magnetic monopoles, especially concerning the current sources.
- **C) The one related to the circulation of the magnetic field and the flux of the electric field:** **Change:** Faraday's Law, relating the circulation of the electric field to the rate of change of the magnetic field, remains:

$$\nabla \times \mathbf{E} = -\frac{\partial \mathbf{B}}{\partial t}$$

This equation does not involve magnetic monopoles directly and thus its form and physical meaning remain unchanged. - **Conclusion:** Faraday's Law ($\nabla \times \mathbf{E}$) is **not different** due to the presence of magnetic monopoles; it continues to describe the fundamental link between time-varying magnetic fields and induced electric fields.

-**D) The ones related to the divergence and the curl of the magnetic field:** - **Change:** We've already established that: - The **divergence of the magnetic field** ($\nabla \cdot \mathbf{B}$) is fundamentally **different** in the presence of magnetic monopoles, changing from $\nabla \cdot \mathbf{B} = 0$ to $\nabla \cdot \mathbf{B} = 4\pi \mathbf{m}$. - The **curl of the magnetic field** ($\nabla \times \mathbf{B}$) is also affected, as Ampère's Law now includes magnetic monopole currents (\mathbf{m}), making it **different** in a universe with magnetic monopoles. - **Conclusion:** Both the divergence ($\nabla \cdot \mathbf{B}$) and the curl ($\nabla \times \mathbf{B}$) of the magnetic field are **altered** in a universe embracing magnetic monopoles, reflecting their direct impact on Gauss's Law for Magnetism and Ampère's Law.

3. **Final Answer Choice**

Given the detailed analysis above, particularly focusing on which Maxwell equations are fundamentally different due to the hypothetical presence of magnetic monopoles:

- **A) The one related to the divergence of the magnetic field $(\nabla \cdot \mathbf{B})^{**}$ is indeed altered, changing from $\nabla \cdot \mathbf{B} = 0$ to $\nabla \cdot \mathbf{B} = 4\pi \mathbf{m}$, directly accounting for magnetic monopole densities. - **B) The ones related to the circulation of the electric field and the divergence of the magnetic field** are different in interpretation and inclusion of magnetic monopole currents, especially in Ampère's Law. - **C) The one related to the circulation of the magnetic field and the flux of the electric field** (Faraday's Law) remains unchanged and is not directly altered by magnetic monopoles. - **D) The ones related to the divergence and the curl of the magnetic field** both show differences: the divergence due to magnetic monopoles and the curl's inclusion of magnetic monopole currents (m) in Ampère's Law.

Conclusion and Final Answer

The **most precise and central equation** that is **fundamentally different** in a universe with magnetic monopoles, directly accounting for the absence of magnetic charge neutrality (i.e., the change from $\nabla \cdot \mathbf{B} = 0$ to $\nabla \cdot \mathbf{B} = 4\pi \mathbf{m}$), is:

This choice highlights the core transformation in Gauss's Law for Magnetism, which is pivotal for understanding how the divergence of the magnetic field ${\bf B}$ is now linked to the distribution of magnetic monopoles (${\bf m}$), fundamentally altering Maxwell's equations in a universe where magnetic monopoles exist.

A is thus the final answer, emphasizing the direct impact of magnetic monopoles on the divergence of the magnetic field, a change that cascades through related equations like Ampère's Law, altering their physical interpretation and current sources (**J** and **m**).

A encapsulates the essence of how Maxwell's equations must adapt in the hypothetical presence of magnetic monopoles, making it the correct choice for understanding the altered behavior of magnetic fields and their sources in such a universe.

[A] is therefore the final, reasoned answer, indicating the equation (Gauss's Law for Magnetism) that is **fundamentally different** in a universe embracing magnetic monopoles, directly due to the non-zero divergence of the magnetic field ($\nabla \cdot \mathbf{B} = 4\pi \mathbf{m}$) in contrast to the monopole-free case ($\nabla \cdot \mathbf{B} = 0$).

A stands as the answer that most accurately reflects the core transformation in Maxwell's equations when magnetic monopoles are introduced, altering the divergence of the magnetic field (B) to account for magnetic charge distributions (m).

Results:

Correct Answer: B Extracted Answer: None

Score: 0.0

Example from Reference-Likert

System Prompt: You are a helpful assistant.

User Prompt: Two quantum states with energies E1 and E2 have a lifetime of 10^{-9} sec and 10^{-8} sec, respectively. We want to clearly distinguish these two energy levels. Which one of the following options could be their energy difference so that they can be clearly resolved?

Please reason step by step, and provide your final answer choice at the end of your response in the format "The final answer is: LETTER" where LETTER is one of A/B/C/D.

Assistant:

To determine which energy difference (E) between the two quantum states with given lifetimes allows for their clear distinction and resolution, we need to carefully analyze the relationship between energy differences, lifetimes, and the quantum mechanical principles governing these states. The key lies in understanding how the energy difference E relates to the lifetimes of the states through the Heisenberg Uncertainty Principle and quantum decay considerations. Let's break this down step by step:

Step 1: Relating Energy Difference to Lifetimes via Quantum Mechanics

The lifetimes of quantum states, especially in the context of energy level transitions and decay, are fundamentally linked to the energy difference between these levels. Specifically, the lifetime τ of a quantum state can be related to the inverse of its effective "decay width" or "resonance width," which itself is influenced by the energy difference E between the states involved in a transition.

Step 2: Heisenberg Uncertainty Principle in Quantum Transitions

The Heisenberg Uncertainty Principle plays a crucial role here, particularly in the context of energy-time uncertainty:

 $\Delta E \cdot \Delta t \ge \frac{\hbar}{2}$

where: - ΔE is the uncertainty in energy (related to the energy difference between states), - Δt is the uncertainty in time (related to the lifetime of the state), - \hbar is the reduced Planck's constant ($\hbar \approx 1.054 \times 10^{-34} \, \mathrm{J} \, \mathrm{s}$).

1350 For our problem, the lifetimes $\tau_1 = 10^{-9}$ sec and $\tau_2 = 10^{-8}$ sec act as the uncertainties in time Δt for the 1351 1352 ### Step 3: Calculating Energy Difference from Lifetimes 1353 To clearly distinguish or resolve the two energy levels (with lifetimes τ_1 and τ_2), the energy difference E 1354 must be significant enough to ensure that the energy uncertainty ΔE is discernible over these lifetimes. 1355 $\Delta E \ge \frac{\hbar}{2 \cdot \Delta t}$ 1356 1357 for both states, where Δt here represents the shorter lifetime τ_{\min} between the two, to ensure the most 1358 stringent resolution condition. Given: $-\tau_1 = 10^{-9} \sec - \tau_2 = 10^{-8} \sec$ 1359 The minimum lifetime τ_{\min} for clear distinction is 10^{-9} sec (the shorter lifetime). 1360 ### Step 4: Applying the Energy-Time Uncertainty to Find E Using $au_{\min}=10^{-9}$ sec in the energy-time uncertainty relation: 1361 1362 $\Delta E \geq \frac{\hbar}{2 \cdot \tau_{\min}}$ 1363 1364 Substituting $\hbar \approx 1.054 \times 10^{-34} \, \mathrm{J} \, \mathrm{s}$ and $\tau_{\mathrm{min}} = 10^{-9} \, \mathrm{sec}$: 1365 $\Delta E \geq \frac{1.054 \times 10^{-34} \, \mathrm{J \ s}}{2 \cdot 10^{-9} \, \mathrm{sec}}$ 1367 $\Delta E \geq \frac{1.054 \times 10^{-34}}{2 \times 10^{-9}}\,\mathrm{J}$ 1369 $\Delta E \ge \frac{1.054}{2} \times 10^{-25} \,\mathrm{J}$ 1370 1371 $\Delta E \ge 0.527 \times 10^{-25} \,\mathrm{J}$ 1372 1373 ### Step 5: Converting Joules to Electronvolts (eV) To compare with the given options in electronvolts (eV), we convert joules to electronvolts using 1 eV =1374 $1.602 \times 10^{-19} \text{ J}$: 1375 $\Delta E \ge 0.527 \times 10^{-25} \,\mathrm{J}$ 1376 $\Delta E \geq \frac{0.527 \times 10^{-25}}{1.602 \times 10^{-19}} \, \mathrm{eV}$ $\Delta E \ge \frac{0.527}{1.602} \times 10^{-6} \,\text{eV}$ 1380 $\Delta E \ge 0.329 \times 10^{-6} \, \mathrm{eV}$ 1381 $\Delta E > 3.29 \times 10^{-7} \, \text{eV}$ 1382 ### Step 6: Comparing with Given Options Now, comparing this derived minimum energy difference $\Delta E \geq 3.29 \times 10^{-7}$ eV with the provided options: 1384 - A) 10^{-9} eV - B) 10^{-11} eV - C) 10^{-4} eV - D) 10^{-8} eV Clearly, the energy difference must be significantly larger than $10^{-9}\,\mathrm{eV}$ and $10^{-11}\,\mathrm{eV}$ to ensure clear dis-1385 tinction, given our derived minimum of 3.29×10^{-7} eV. 1386 Step 7: Final Selection Based on Clarity and Resolution 1387 Given the requirement for clear distinction and resolution between the two quantum states with lifetimes 1388 10^{-9} sec and 10^{-8} sec: 1389 - Option C) 10⁻⁴ eV stands out as the most appropriate choice, as it far exceeds the derived minimum 1390 energy difference of 3.29×10^{-7} eV, ensuring a substantial energy gap that allows for clear resolution and 1391 distinction between the two quantum states. ### Conclusion and Final Answer 1392 Thus, the energy difference E that allows for the clear distinction and resolution of the two quantum states 1393 (with lifetimes 10^{-9} sec and 10^{-8} sec) is best captured by option C) 10^{-4} eV, due to its substantial magni-1394 tude compared to the derived minimum energy uncertainty. 1395 C 1396 This final answer choice, C, ensures that the energy difference E is sufficiently large to clearly resolve 1398 the two quantum states with their given lifetimes, adhering to the principles of quantum mechanics and the 1399 Heisenberg Uncertainty Principle for energy-time uncertainties. 1400 **Results:** 1401 Correct Answer: C 1402 Extracted Answer: C

1403

Score: 1.0

Example from Qwen2.5-7B

System Prompt: You are a helpful assistant.

User Prompt: In a parallel universe where a magnet can have an isolated North or South pole, Maxwell's equations look different. But, specifically, which of those equations are different?

A) The one related to the divergence of the magnetic field. B) The ones related to the circulation of the electric field and the divergence of the magnetic field. C) The one related to the circulation of the magnetic field and the flux of the electric field. D) The ones related to the divergence and the curl of the magnetic field. Please reason step by step, and provide your final answer choice at the end of your response in the format "The final answer is: LETTER" where LETTER is one of A/B/C/D.

Assistant:

To determine which Maxwell's equations would be different in a universe where a magnet can have an isolated North or South pole, we need to consider the implications of such a scenario on the fundamental properties of magnetic fields.

- 1. **Divergence of the Magnetic Field**: In our universe, the divergence of the magnetic field is always zero, meaning there are no magnetic monopoles. If a magnet could have an isolated North or South pole, this would imply the existence of magnetic monopoles, which would change the divergence of the magnetic field to be non-zero at the location of the monopole. Therefore, the equation related to the divergence of the magnetic field would be different.
- 2. **Circulation of the Electric Field and Divergence of the Magnetic Field**: These equations describe how electric fields are generated by changing magnetic fields and vice versa. The presence of magnetic monopoles does not directly affect these equations, as they are primarily concerned with the relationship between electric and magnetic fields over time and space.
- 3. **Circulation of the Magnetic Field and Flux of the Electric Field**: These equations relate to Faraday's law, which describes how changing magnetic fields induce electric fields. Again, the presence of magnetic monopoles does not directly impact these equations, as they focus on the interaction between magnetic fields and electric fields.
- 4. **Divergence and Curl of the Magnetic Field**: The divergence of the magnetic field is already discussed above. The curl of the magnetic field is related to the circulation of the electric field and is not directly affected by the existence of magnetic monopoles.

Based on this reasoning, the equation that would be different in a universe where a magnet can have an isolated North or South pole is the one related to the divergence of the magnetic field.

The final answer is: |A|

Results:

Correct Answer: B Extracted Answer: A

Score: 0.0

Example from Qwen2.5-7B

System Prompt: You are a helpful assistant.

User Prompt: Two quantum states with energies E1 and E2 have a lifetime of 10^{-9} sec and 10^{-8} sec, respectively. We want to clearly distinguish these two energy levels. Which one of the following options could be their energy difference so that they can be clearly resolved?

```
A) 10^{-9} \text{ eV}
```

B) 10^{-11} eV

C) 10^{-4} eV

 $\stackrel{\frown}{\rm D}) 10^{-8} \, {\rm eV}$

Please reason step by step, and provide your final answer choice at the end of your response in the format "The final answer is: LETTER" where LETTER is one of A/B/C/D.

Assistant

To determine which energy difference allows us to clearly distinguish the two energy levels, we need to consider the uncertainty principle. The uncertainty in energy (E) multiplied by the uncertainty in time (t) should be greater than or equal to Planck's constant (h) divided by 4. In this case, the uncertainty in time is the difference between the lifetimes of the two states.

E * t h / (4)

```
1458
           Given that the lifetimes are 10^-9secand10^-8sec, the uncertainty in time is:
1459
           t = 10^{-}8 - 10^{-}9 = 9 * 10^{-}10sec
1460
           Now, we can plug this value into the equation and solve for E:
1461
           E * 9 * 10^- 10h/(4)
1462
           E (h/(4))/(9*10^{-10})
           Using the value of Planck's constant (h 6.626 * 10^{-}34Js), we get:
1463
           E (6.626 * 10^{-}34)/(4 * 3.14159 * 9 * 10^{-}10)
1464
           E 5.8 * 10^-25J
1465
          Now, we need to convert this energy difference to electron volts (eV) by dividing by the charge of an electron
1466
           (e 1.602 * 10^-19C):
           E (5.8 * 10^{-}25)/(1.602 * 10^{-}19)
1467
           E 3.6 * 10^{-}6eV
1468
          Comparing this value to the given options, we can see that the closest option is:
1469
          The final answer is: C
1470
1471
           Results:
1472
           Correct Answer: C
1473
           Extracted Answer: C
           Score: 1.0
1474
1476
1477
1478
1479
1480
1481
1482
1483
1484
1485
1486
1487
1488
1489
1490
1491
1492
1493
1494
1495
1496
1497
1498
1499
1500
```