

000 FROM <ANSWER> TO <THINK>: MULTIDIMEN- 001 SIONAL SUPERVISION OF REASONING PROCESS FOR 002 LLM OPTIMIZATION 003 004

006 **Anonymous authors**

007 Paper under double-blind review

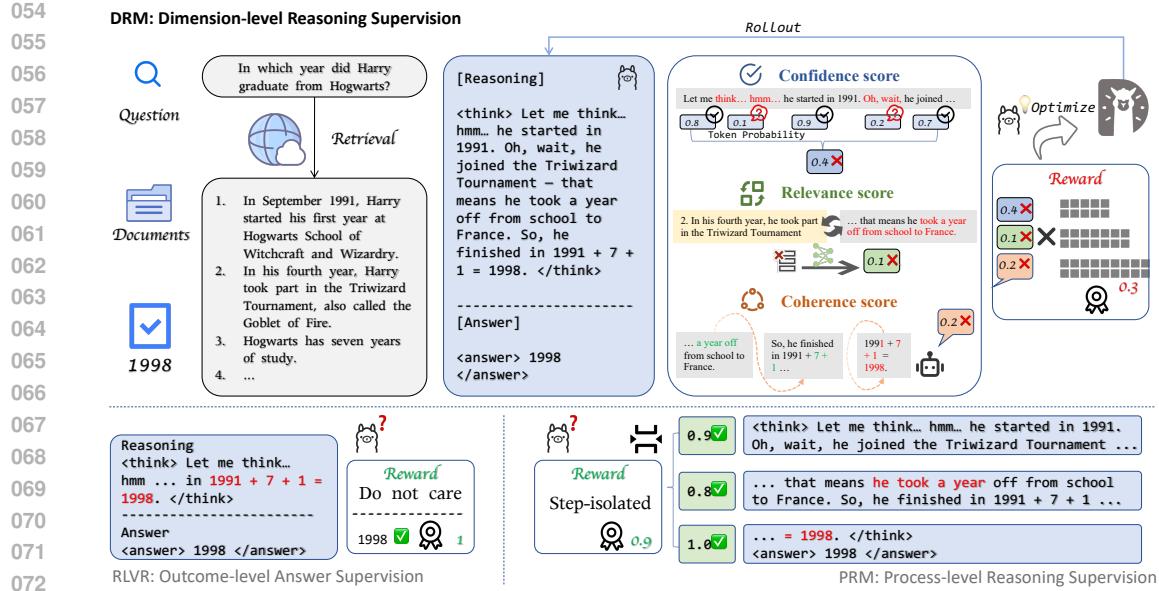
010 ABSTRACT

013 Large language models (LLMs) can develop strong reasoning ability when trained
014 appropriately. Existing approaches are broadly categorized into outcome-level
015 answer supervision and process-level reasoning supervision. However, the for-
016 mer provides only sparse binary feedback and overlooks intermediate step qual-
017 ity, while the latter scores individual steps but requires task-specific segmentation.
018 To this end, we propose a novel framework that assesses the quality of reasoning
019 process along three dimensions: **Confidence** for uncertainty calibration, **Rele-**
020 **vance** for semantic alignment and **Coherence** for logical consistency. Together,
021 these dimensions capture aspects beyond final answer correctness and enable in-
022 terpretable assessment without requiring ground truth answers. Our framework
023 serves as a **Dimension-level Reward Model (DRM)** that assigns scores to reason-
024 ing processes and provides supervision signals for both off-policy (e.g., DPO) and
025 on-policy (e.g., GRPO) optimization. Experimental results show that DRM pro-
026 vides effective supervision signals, guides the optimization of LLMs and enhances
027 their reasoning ability. In particular, DRM-supervised training achieves consistent
028 gains on both in-distribution and out-of-distribution open-domain tasks, includ-
029 ing mathematics, question answering, code execution and puzzles. Our findings
030 demonstrate that multidimensional supervision of reasoning process can improve
031 the generalized reasoning ability of LLMs beyond the training distribution.

032 1 INTRODUCTION

034 Enhancing the reasoning ability of Large Language Models (LLMs) to perform complex and multi-
035 step reasoning remains a central challenge in their development (Zhang et al., 2025b; Xu et al.,
036 2025). The dominant paradigm for enhancement relies on Reinforcement Learning with Verifiable
037 Rewards (RLVR) (Shao et al., 2024; Yang et al., 2024; Luo et al., 2024). RLVR provides supervision
038 at the outcome level, assigning a positive reward only if the final answer is correct. However, this re-
039 ward mechanism has fundamental limitations. First, answer supervision overlooks the quality of the
040 reasoning process (Yu et al., 2025a). This often leads to rewarding models for arriving at a *correct*
041 *answer with flawed reasoning* while penalizing sound logic that contains a minor final error (Xie
042 et al., 2025). Second, we observed that rewards in RLVR can become nearly constant when the
043 model is either too powerful or too weak on the training set, thereby offering limited guidance for
044 optimization (Cui et al., 2025). Process-level Reward Models (PRMs) are designed to address these
045 limitations by supervising intermediate steps (Cheng et al., 2025; Zhang et al., 2025a; Zou et al.,
046 2025). While promising, PRMs introduce their own challenges. Their process-level supervision
047 requires the reasoning process to be segmented into individual steps (Xiong et al., 2025; Zou et al.,
048 2025). This segmentation is often learned in a task-specific manner, which may hinder generaliza-
049 tion to open-domain tasks with ambiguous or overlapping steps (Xiong et al., 2025). Furthermore,
050 unlike the transparent binary signal of RLVR, PRMs often function as black-box evaluators, making
it difficult to diagnose or trust their scoring mechanism (Christiano et al., 2023).

051 To overcome these limitations, we propose a new supervision framework grounded in the key char-
052 acteristics of a high-quality reasoning process. Prior work shows that unfaithful content in reasoning
053 process can hinder correct answers (Zhang et al., 2025b). To detect such content, our framework
performs assessment along three complementary dimensions: (1) **Confidence**, measures whether



073
074 Figure 1: An overview of our multidimensional reasoning supervision framework, illustrated on a RAG task. RLVR regards a *correct answer with flawed reasoning* as a positive sample since it
075 focuses solely on the answer. PRMs also misclassify it because process-level supervision ignores
076 errors across steps when each individual step is correct. **DRM** performs dimension-level supervi-
077 sion, detects reasoning flaws, and assigns a reward that reflects the real quality of reasoning process,
078 facilitating further optimization.

080
081 the reasoning remains faithful to the question and supporting context, directly counters the *flawed reasoning* issue where models hallucinate or deviate; (2) **Relevance**, evaluates the semantic rela-
082 tionship and entailment between the reasoning process and the question, the supporting context and
083 the final answer, enabling the detection of deviations from the given information; and (3) **Coher-
084 ence**, penalizes self-contradictory statements by the logical consistency of the reasoning process.
085 Figure 1 illustrates how our framework as-
086 sesses the quality of the reasoning process as
087 a **Dimension-level Reward Model (DRM)** and
088 addresses the limitations of both RLVR and
089 PRMs. Table 1 summarizes the key properties
090 of the three supervision approaches. By provid-
091 ing a dense, reasoning-aware reward signal with-
092 out requiring task-specific ground truth answers,
093 DRM overcomes the key limitations of RLVR.
094 Simultaneously, it avoids the task-specific seg-
095 mentation required by PRMs and offers superior interpretability by scoring reasoning along explicit,
096 diagnosable dimensions.

097 Experimental results on multiple challenging open-domain benchmarks demonstrate the effective-
098 ness of DRM-based supervision in both off-policy selection and on-policy training paradigms. Our
099 results show that DRM-supervised models perform competitively on both in-distribution and out-
100 of-distribution tasks, indicating stronger generalization than answer-supervised counterparts. For
101 LLAMA-3.1-8B-INSTRUCT (Grattafiori et al., 2024), our method achieves performance gains on
102 MATH500 (+8.8, mathematics) (Cobbe et al., 2021a), 2WIKI-RAG (+8.7, multi-hop QA) (Ho
103 et al., 2020) and CRUXEVAL (+7.1, code execution) (Gu et al., 2024). This improvement trend
104 is consistently observed across different models, which unequivocally demonstrates the superiority
105 and generality of DRM supervision. Qualitative analysis and case studies show that DRM mitigates
106 the *correct answer with flawed reasoning* issue common in answer supervision. Our results indi-
107 cate that multidimensional reasoning supervision enhances the reasoning ability of LLMs and their
108 performance on out-of-distribution tasks.

Table 1: Comparison of supervision approaches.

Property	RLVR	PRM	DRM
Supervision level	Outcome	Process	Dimension
Supervision target	Answer	Reasoning	Reasoning
Dense signal	✗	✓	✓
Generalization	✓	✗	✓
Interpretability	✓	✗	✓
Ground truth free	✗	✓	✓

108 Table 2: Reasoning assessment dimensions, following the (Q, D, R, A) quadruple format.
109

110 Dimension	111 Description	112 Implementation
113 Confidence score ^{Conf}	114 Self-assessed certainty of generated R and A from intrinsic signals.	115 score ^{Conf} _{R} = $\frac{1}{ R } \sum \log p$, for all tokens in R . score ^{Conf} _{A} = $\sum \log p$, for all tokens in A . score ^{Conf} = score ^{Conf} _{R} + score ^{Conf} _{A} .
116 Relevance score ^{Rel}	117 Evaluates whether R is contextually appropriate and semantically aligned with Q, D and A .	118 $R \leftarrow Q$: Measured by NLI entailment. $R \leftrightarrow D$: Measured by semantic relevance. $R \rightarrow A$: Measured by NLI entailment.
119 Coherence score ^{Coh}	120 Evaluates logical consistency, fluency and overall quality of R .	121 Evaluated by an external ORM.

122

2 METHODOLOGY: MULTIDIMENSIONAL REASONING SUPERVISION

123
124 **Task Definition.** Formally, let I denote the user input and O the model output. We decompose O
125 into a reasoning process R and an answer A . In open-domain scenarios, I often contains more than
126 just the question Q . For example, in Retrieval-Augmented Generation (RAG) tasks, I additionally
127 includes retrieved documents, while in preference tasks, I may consist of two candidate responses
128 for the model to compare. Let D denote the additional information accompanying Q and we can
129 decompose I into Q and D . Consequently, the input–output structure of the model can be denoted
130 by a quadruple: (Q, D, R, A) . In most tasks, the performance of the model is evaluated primarily
131 based on the quality of A .132 Prior work shows that LLMs sometimes generate unsupported statements during reasoning, which
133 can hinder the production of correct answers (Zhang et al., 2025b; Xu et al., 2025). To address this
134 issue, models are expected to produce faithful reasoning that avoids unsupported claims. In particu-
135 lar, they should produce decisive output, especially for the final answer. Furthermore, the reasoning
136 process should be grounded in the provided input and exhibit internal consistency throughout. These
137 properties support both the production of correct answers and the interpretability of reasoning pro-
138 cess. We categorize these properties into three dimensions that a high-quality reasoning process
139 should satisfy: **Confidence**, **Relevance** and **Coherence**. Table 2 summarizes their definitions and
140 implementation and the rationale for each is discussed in the following.141 **Confidence.** This dimension evaluates whether the models are certain about their output. Inspired
142 by prior work on self-confidence evaluation in reasoning models, we compute the average log-
143 probability of tokens in R (Leang et al., 2025) to avoid penalizing exploratory reasoning processes.
144 For A , we compute the sum of log-probability instead to encourage decisive and confident outputs.
145 The final confidence score is calculated as the sum of these two components.146 **Relevance.** This dimension assesses whether R maintains necessary relationships with other
147 components Q, D and A : (1) $Q \rightarrow R$ should hold via Natural Language Inference (NLI) entailment,
148 ensuring R contributes to answering Q ; (2) $R \leftrightarrow D$ should exhibit high semantic relevance, ensuring
149 R is grounded in the additional information D ; and (3) $R \rightarrow A$ should also hold via NLI entailment,
150 ensuring R logically leads to A . Specifically, we compute the relevance score by framing it as a
151 ranking task: we rank the reasoning process using three distinct metrics, each corresponding to one
152 of the relationships defined earlier, and then combine these scores to obtain the final score.153 **Coherence.** This dimension evaluates the text quality of the reasoning process, with attention
154 to coherence and logical consistency. We treat R as the output of a text generation task with the
155 input of Q, D . To assess its logical consistency, fluency, and overall textual quality, we use an
156 external Outcome-level Reward Model (ORM) in the text-quality evaluation. This captures another
157 dimension of reasoning quality that is not directly reflected in confidence or relevance.158 Overall, by jointly evaluating the reasoning process along **Confidence**, **Relevance** and **Coherence**,
159 our framework explicitly decomposes assessment into complementary dimensions. As illustrated in
160 Figure 1, DRM assesses reasoning quality along three distinct dimensions with each grounded in
measurable scores. We compute the DRM reward by a weighted sum of the dimensional scores:

161
$$R_i^{DRM} = \text{score}_i = \sum_D w^D \tilde{\text{score}}_i^D, \quad D \in \{\text{Conf}, \text{Rel}, \text{Coh}\},$$

162 where $\widetilde{\text{score}}_i^D$ is the component score i^D after being individually normalized within its group to mitigate scale differences. This produces a dense reward that serves as a direct supervision signal. The weights are determined via a grid search on the validation set. This design inherently avoids the binary sparse reward issue of RLVR and reflects the quality of the reasoning process. DRM replaces stepwise scoring with dimension-wise assessment and eliminates the need for task-specific step segmentation in PRMs. Owing to its dimensional nature, DRM inherently provides more interpretable feedback. Moreover, it can distinguish among multiple reasoning processes by their quality, regardless of answer correctness. As DRM addresses the evaluation limitations of RLVR and PRM, we investigate whether its reward can serve as an effective supervision signal for LLM optimization. In off-policy optimization, training sets are constructed under the guidance of a supervision signal. R_i^{DRM} can serve this role by capturing the reasoning quality of each sample, thereby facilitating training set construction. We adopt DPO, and its optimization objective is formulated as follows:

$$\mathcal{L}_{\text{DPO}}(\theta) = -\mathbb{E}_{(I, O^+, O^-)} \left[\log \sigma \left(\beta \log \frac{\pi_\theta(O^+ | I)}{\pi_{\text{ref}}(O^+ | I)} - \beta \log \frac{\pi_\theta(O^- | I)}{\pi_{\text{ref}}(O^- | I)} \right) \right],$$

$$O^+ = \arg \max_{o \in O} R_o^{DRM}, \quad O^- = \arg \min_{o \in O} R_o^{DRM},$$

174 where $\sigma(\cdot)$ is the sigmoid function and $\beta > 0$ controls the sharpness of preference. In on-policy 175 optimization, DRM can serve as a standalone supervision reward signal, or be integrated with other 176 supervision signals. Specifically, we compute an additional DRM advantage $\hat{A}_{i,t}^{DRM}$ from R_i^{DRM} , 177 which denotes the DRM reward for sample i . We then add this DRM advantage to the native GRPO 178 advantage $\hat{A}_{i,t}$ obtained from RLVR rewards, yielding our optimization objective (for mathematical 179 details, please refer to Appendix B.2):

$$\mathcal{J}_{\text{GRPO}}(\theta) = \mathbb{E}_{q, \{o_i\}} \frac{1}{G} \sum_{i=1}^G \frac{1}{|o_i|} \sum_{t=1}^{|o_i|} \left\{ \min [r_{i,t}(\theta) A_{i,t}, \text{clip}(r_{i,t}(\theta), 1 - \varepsilon, 1 + \varepsilon) A_{i,t}] \right. \\ \left. - \beta \mathbb{D}_{\text{KL}} [\pi_\theta \| \pi_{\text{ref}}] \right\}, \quad A_{i,t} = \begin{cases} \hat{A}_{i,t}, & \text{RLVR,} \\ \hat{A}_{i,t}^{DRM}, & \text{DRM,} \\ \hat{A}_{i,t} + \hat{A}_{i,t}^{DRM}, & \text{Combination of RLVR and DRM,} \end{cases}$$

194 where $r_{i,t}(\theta) = \frac{\pi_\theta(o_{i,t} | q, o_{i,<t})}{\pi_{\theta_{\text{old}}}(o_{i,t} | q, o_{i,<t})}$ is the token-level probability ratio and β controls the KL penalty 195 strength with respect to a reference policy π_{ref} . DRM can be employed either as a standalone signal 196 or integrated with the RLVR supervision signal.

3 EXPERIMENTS

201 Following a rigorous experimental paradigm, we formulate a set of research questions (RQs) to 202 evaluate whether DRM supervision can improve the model's reasoning ability. The empirical results 203 presented in this section affirmatively answer all of the following research questions.

204 **RQ1:** *Can assessment on reasoning process reliably determine the final answer correctness?*
 205 **RQ2:** *Can the DRM reward signal be learned and used by models to improve reasoning ability?*
 206 **RQ3:** *Can DRM supervision better guide training and outperform RLVR?*
 207 **RQ4:** *Can combining RLVR supervision with DRM supervision lead to further improvements?*

3.1 EXPERIMENTAL SETUP

212 **Models.** We evaluate our method on three representative models: a model lacking inherent 213 reasoning ability **LLAMA-3.1-8B-INSTRUCT** (Grattafiori et al., 2024), a reasoning model **R1- 214 DISTIL-LLAMA8B** (DeepSeek-AI et al., 2025), and a hybrid reasoning model **QWEN3-8B** (Yang 215 et al., 2025). We employ QWEN3-8B-RERANKER (Zhang et al.) as the relevance judge and LLAMA- 216 3.3-NEMOTRON-70B-REWARD-MULTILINGUAL (Wang et al.) as the coherence judge.

216 Table 3: Answer correctness (%) of DRM construction approaches on RewardBench2. Native means
 217 the performance of the backbone models. (0.1,0.2,0.7) means weights for Confidence, Relevance
 218 and Coherence are 0.1, 0.2, 0.7, respectively. LTR denotes the use of a Learning-to-Rank model
 219 with learnable weights for integration. The highest result in each row is in **bold**.
 220

221 Model	222 Native	223 Confidence	224 Relevance	225 Coherence	226 Weighted 227 Equally	228 Weighted 229 (0.1,0.2,0.7)	230 LTR
223 LLaMA3.1-8B-Instruct	224 67.17	225 65.44	226 72.32	227 78.55	228 77.45	229 78.57	230 79.13
223 R1-Distil-Llama8B	224 63.46	225 63.10	226 66.76	227 76.35	228 75.11	229 76.16	230 75.18
223 Qwen3-8B	224 84.87	225 83.20	226 85.10	227 85.54	228 85.01	229 85.65	230 85.88

226
 227 **Datasets.** We evaluate our method on a diverse set of open-domain tasks, including four **Code**
 228 benchmarks, two **Preference** benchmarks, four **Math** benchmarks, two **Scientific QA** benchmarks,
 229 three **Logical Reasoning** benchmarks and two **Question Answering** benchmarks along
 230 with their RAG variants provided by FlashRAG (Jin et al., 2024). For math tasks, we use **MATH-
 231 VERIFY** (Kydlíček, 2024) for automatic solution verification and **exact match** for all other tasks.¹
 232

233 3.2 EVALUATING WHETHER DRM GUIDES CORRECT ANSWERS

234 To address **RQ1**, we validate the effectiveness of DRM using a Best-of-N (BoN) selection setup. The
 235 underlying hypothesis is that a high-quality reasoning process assessed by our multi-dimensional
 236 reward serves as a reliable proxy for answer correctness. Specifically, for each test instance, we
 237 sample multiple candidate reasoning paths from the model and select the one with the highest DRM
 238 reward. We then evaluate whether this selection mechanism yields higher answer accuracy compared
 239 to three types of baselines: a baseline obtained via uniform sampling of reasoning processes, which
 240 reflects the model’s native performance in the absence of explicit supervision signals; baselines using
 241 each individual DRM dimension (**Confidence**, **Relevance**, or **Coherence**) in isolation, which allows
 242 us to assess the contribution of each signal separately; and a baseline where these three dimensions
 243 are integrated with equal weights. Furthermore, we also compare fixed weighting schemes against
 244 learnable weights. We employ a Learning-to-Rank (LTR) approach based on LambdaRank (Burges
 245 et al., 2007; Burges, 2010), training the model to optimize the combination of dimensional scores to
 246 maximize the probability of correctness.

247 As shown in Table 3, DRM consistently achieves higher accuracy than the backbone models. While
 248 using the **Confidence** score alone slightly reduces accuracy, combining it with **Relevance** and **Co-
 249 herence** improves performance, indicating that these dimensions capture complementary aspects of
 250 reasoning quality. Regarding integration mechanisms, the combined approach consistently outper-
 251 forms both individual metrics and native backbone performance, regardless of whether the integra-
 252 tion employs equal weighting, grid-search fixed weights or a learnable mechanism. This stability
 253 is observed across diverse backbones and is further validated on a distinct data distribution (Hot-
 254 potQA with RAG) in Table 11. Given that the performance gap between fixed weights and the more
 255 complex LTR approach is marginal, we determine the combination weights via grid search on the
 256 validation set and fix them for all subsequent experiments. This choice prioritizes simplicity and
 257 robustness, eliminating the need for additional training to learn parameters. Overall, the results of
 258 our extensive experiments demonstrate that DRM maintains robustness across different backbone
 259 models, integration methods and training data distributions.

260 3.3 ASSESSING THE EFFECTIVENESS OF DRM SUPERVISION

261 This section focuses on **RQ2** and **RQ3**. We conduct off-policy reinforcement learning using DPO
 262 with Supervised Fine-Tuning (SFT) loss (for mathematical details, please refer to Appendix B.1).
 263 We construct separate training sets based on different supervision signals. Specifically, DRM
 264 rewards serve as reasoning supervision signals, guiding the selection of samples with higher reason-
 265 ing quality, while RLVR rewards serve as answer supervision signals, selecting samples based on
 266 answer correctness. For each instance in RewardBench2, we prompt the model to generate 20 sam-
 267

268 ¹The main paper only reports results on RewardBench2; results for HotpotQA with RAG are provided in
 269 Appendix G.

270 ples containing step-by-step reasoning and final answers. These samples are scored and selected
 271 according to the respective supervision signal to form preference pairs, as described below.
 272

273 **Training Set Construction.**

274 Let x denote a sample from set X , where all samples in X are generated from the same instance.
 275 Each sample is associated with a correctness label $\text{answer}_x \in \{\text{True}, \text{False}\}$ and a reasoning
 276 quality score score_x . The positive set X^+ and negative set X^- are defined according to a **SUBSET**
 277 rule and preference pairs are selected according to a **SUPERVISION** method. Once these two
 278 components are specified, the resulting training set is uniquely determined.

279 **SUBSET:**

280 **ANY:** $X^+ = X^- = X$.

281 **T+T:** $X^+ = X^- = \{x \mid \text{answer}_x = \text{True}, x \in X\}$.

282 **T+F:** $X^+ = \{x \mid \text{answer}_x = \text{True}, x \in X\}$, $X^- = \{x \mid \text{answer}_x = \text{False}, x \in X\}$.

283 **F+F:** $X^+ = X^- = \{x \mid \text{answer}_x = \text{False}, x \in X\}$.

284 **SUPERVISION**

285 **DRM:** $\{(x^+, x^-) \mid x^+ = \arg \max_{x \in X} \text{score}_x, x^- = \arg \min_{x \in X} \text{score}_x\}$

286 **RLVR:** $\{(x^+, x^-) \mid x^+ = \text{random}(X^+), x^- = \text{random}(X^-)\}$

287 Let **SUPERVISION@SUBSET** denote a training set construction method. For example,
 288 **DRM@T+F** indicates that we select a sample with the highest DRM reward and correct
 289 answer and pair it with a sample with the lowest DRM reward and wrong answer. It is clear
 290 that **DRM@ANY** refers to the training set constructed with DRM supervision. In contrast,
 291 **RLVR@T+F** refers to the training set constructed with answer supervision, under the RLVR
 292 assumption that samples with the same answer are considered equivalent.

293 We construct separate training sets and train models on each set independently. The full training
 294 details are provided in Appendix F.3. As shown in Table 4, DRM-supervised training consistently
 295 outperforms RLVR-supervised training, providing evidence in support of both research questions.

296 **RQ2.** To assess whether DRM reward signals can be effectively learned and used to improve
 297 reasoning ability, we compare **NATIVE** and **DRM@ANY**. Additionally, we include **RLVR@ANY**
 298 as a control group, in which the training set was constructed randomly. In the **DRM@ANY** setting,
 299 the training set is constructed entirely based on DRM reward signals, without incorporating
 300 any information about answer correctness. Table 4 shows that **DRM@ANY** achieves higher scores
 301 than all other settings, with substantial improvements across all evaluated datasets. The strong
 302 performance on out-of-distribution tasks suggests that the model generalizes well beyond the training
 303 distribution. The results indicate that the proposed DRM supervision can be effectively learned even
 304 without answer supervision, i.e., without access to the ground truth answers.

305 **RQ3.** We compare DRM and RLVR across two key aspects to assess their relative effectiveness:
 306 **Performance gain:** To evaluate the effectiveness of DRM, we compare **RLVR@T+F** with
 307 **DRM@ANY** (see Table 4). This comparison examines whether explicit supervision of reasoning
 308 achieves better performance than supervising only the answer. In this setting, **DRM@ANY**
 309 consistently achieves higher performance than **RLVR@T+F**, indicating that training with DRM
 310 supervision consistently outperforms RLVR supervision.

311 **Overcoming limitations:** We compare **RLVR@T+T** with **DRM@T+T** and **RLVR@F+F** with
 312 **DRM@F+F** to test whether DRM can still provide supervision when all answers have identical
 313 correctness labels, where RLVR cannot produce a preference signal. Results show that DRM can
 314 distinguish reasoning quality in such case, demonstrating its ability to generate informative supervision
 315 and to enhance the model’s ability to handle a broader range of scenarios.

316 Furthermore, we conduct off-policy training and compare it against the baselines as shown in Ta-
 317 ble 5. We evaluate our model against three strong baselines: (1) a model trained on the **ANY** sub-
 318 set with reasoning supervision signals from SKYWORK-REWARD-V2-LLAMA-3.1-8B, a powerful
 319 ORM, (2) RLPR (Yu et al., 2025b) and (3) KLEAR (Su et al., 2025). Both RLPR and KLEAR are
 320 reasoning-enhanced models trained using the same backbone architecture as their counterparts in our
 321 experiments. This setup allows us to examine whether our DRM provides more effective and
 322 generalizable supervision than existing reasoning-supervision approaches. We also examines whether
 323 DRM-supervised models can outperform models optimized with other methods. Across most down-
 324 stream open-domain tasks, DRM outperforms all three baselines. In particular, it surpasses RLPR
 325 and KLEAR under the same backbone, demonstrating its effectiveness. It also exceeds the perfor-

324
 325 Table 4: Results of controlled comparisons for RQ2 and RQ3. We use LLAMA3.1-8B-INSTRUCT
 326 as the base model. Results for other models, which exhibit the same trend, are provided in Ap-
 327 pendix G.2. As described in Section 3.1, we use MATH-VERIFY as the evaluation metric for math
 328 tasks and EM for all other tasks, respectively. All models are trained for the same number of steps
 to ensure a fair comparison. For each row within a comparison, the highest score is in **bold**.

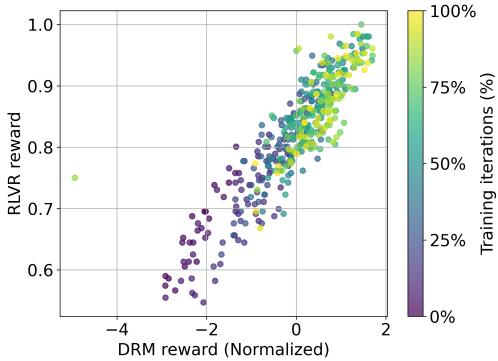
Task Domain	Dataset	For RQ2, RQ3.1				For RQ3.2			
		Native	RLVR @ ANY	RLVR @ T+F	DRM @ ANY	RLVR @ T+T	DRM @ T+T	RLVR @ F+F	DRM @ F+F
Code	CodeMMLU	58.8	58.8	59.5	59.9	58.9	59.6	59.6	61.3
	CodeScope	34.8	35.4	37.4	41.1	36.2	41.0	36.6	40.0
	Cruxeval	50.4	53.5	52.6	57.5	53.6	56.6	53.9	55.9
	Execution-v2	38.2	40.9	43.2	45.3	39.2	45.5	40.3	46.8
Preference	RM-Bench	56.4	59.3	59.2	61.0	60.0	60.3	59.7	61.9
	UltraFeedback	66.6	65.6	65.4	69.9	66.4	67.7	64.5	68.8
Math	AIME24	4.7	4.7	4.0	6.0	4.7	7.3	4.7	4.0
	AMC23	22.5	23.5	23.5	29.5	23.0	25.5	22.0	26.5
	GSM8K	88.8	89.0	89.5	91.8	90.2	91.7	88.7	91.7
	Math500	39.6	41.4	43.4	48.4	42.0	46.6	40.0	48.4
Scientific QA	MMLU-Pro	41.9	45.3	46.4	48.7	45.7	48.4	46.6	49.0
	GPQA	31.3	28.8	32.8	35.9	29.8	30.3	29.8	35.4
Reasoning	MuSR	48.3	49.5	49.7	51.7	48.3	53.3	49.7	51.6
	DROP	56.9	61.0	62.9	63.6	60.0	64.4	58.5	65.1
	QASC	84.4	84.0	84.2	87.2	83.8	87.8	83.4	86.2
QA	2wiki	33.8	33.2	34.6	35.6	32.3	32.7	30.7	33.4
	HotpotQA	29.3	29.9	30.1	31.8	29.3	30.1	29.1	29.7
QA-RAG	2wiki_RAG	31.2	32.1	35.8	39.9	36.6	41.4	32.1	43.3
	HotpotQA_RAG	28.3	28.3	32.3	34.5	29.3	32.3	28.5	33.8

351
 352 mance of the model trained with SKYWORK supervision, indicating that DRM consistently achieves
 353 stronger and more generalizable reasoning ability. The improvements are consistent across various
 354 architectures and tasks, suggesting that DRM is an architecture-agnostic approach that generalizes
 355 well. Notably, our training relies solely on preference data from RewardBench2, the same type
 356 of data used for training reward models (Zhang et al., 2025a; Zhong et al., 2025), without access
 357 to ground truth answers or task-specific finetuning. This highlights the data efficiency of our ap-
 358 proach as a single source of preference data leads to broad improvements across open-domain tasks.

3.4 ENHANCING RLVR WITH DRM

361
 362 This section addresses **RQ4**. We conduct on-
 363 policy GRPO training on three advantage con-
 364 figurations: answer supervision only, reason-
 365 ing supervision only and their combination.
 366 This setup directly tests whether DRM supervi-
 367 sion and integrating DRM rewards into RLVR
 368 achieve further gains. The comparison between
 369 RLVR and DRM also examines whether the
 370 trend observed in off-policy training remains
 371 consistent in on-policy stages. GRPO training
 372 details are provided in Appendix F.4.

373 Across most model backbones and represen-
 374 tative benchmarks on open-domain tasks, the
 375 combined approach performs as well as or bet-
 376 ter than the best single supervision approach,
 377 as shown in Table 6. This trend is also consis-
 378 tently observed in the off-policy setting. The
 379 combination also outperforms RLVR, indicat-



380
 381
 382 Figure 2: The relationship between RLVR re-
 383 wards and DRM rewards in R1-Distil-Llama8B
 384 Combination training. Each data point repre-
 385 sents a single training batch. Note that DRM re-
 386 wards are Z-score normalized for better visualiza-
 387 tion.

Table 5: Results of off-policy DPO with SFT loss training. **RLPR** and **KLEAR** are baseline models that share the same backbone architectures as their respective counterparts. **SKYWORK** indicates that the model’s training set is constructed using SKYWORK reward model. **DRM** represents **DRM@ANY**. For each row within a model group, the highest score is in **bold**.

Task Domain	Dataset	LLaMA3.1-8B-Instruct				R1-Distil-Llama8B				Qwen3-8B			
		Native	RLPR	SKYWORK	DRM	Native	SKYWORK	DRM	Native	KLEAR	SKYWORK	DRM	
Code	CodeMMLU	58.8	58.0	57.6	59.9	59.7	62.9	66.3	77.9	77.4	79.3	80.3	
	CodeScope	34.8	38.7	39.3	41.1	67.4	68.2	70.2	86.5	88.1	86.2	87.4	
	Cruxeval	50.4	53.6	53.6	57.5	71.9	77.0	77.2	91.6	87.2	91.9	93.0	
	Execution-v2	38.2	44.7	42.8	45.3	80.8	82.0	86.0	98.5	95.2	97.9	99.0	
Preference	RM-Bench	56.4	60.2	59.8	61.0	71.9	73.4	74.6	85.4	83.7	85.1	85.6	
	UltraFeedback	66.6	68.5	67.0	69.9	65.2	66.5	66.8	71.3	68.1	72.2	73.2	
Math	AIME24	4.7	6.0	4.0	6.0	28.7	26.7	33.3	38.0	40.0	38.7	44.7	
	AMC23	22.5	26.0	25.5	29.5	70.5	74.5	75.5	72.0	75.0	76.0	79.0	
	GSM8K	88.8	90.0	89.8	91.8	66.7	73.7	69.2	95.6	93.8	95.8	96.1	
	Math500	39.6	47.2	42.6	48.4	62.6	65.6	63.2	73.2	68.2	72.6	75.6	
Scientific QA	MMLU-Pro	41.9	36.3	46.7	48.7	51.5	52.8	54.7	65.3	67.1	70.0	71.4	
	GPQA	31.3	30.8	33.3	35.9	39.9	37.4	44.9	48.0	55.6	52.5	58.1	
Reasoning	MuSR	48.3	48.7	49.7	51.7	52.6	52.8	54.1	63.5	50.8	63.5	65.5	
	DROP	56.9	45.4	63.0	63.6	50.8	54.5	50.2	74.7	68.8	74.2	74.9	
	QASC	84.4	87.0	87.1	87.2	82.1	82.5	84.1	94.1	93.3	93.7	94.2	
QA	2wiki	33.8	32.1	32.4	35.6	26.2	29.3	31.6	39.8	35.9	40.0	42.2	
	HotpotQA	29.3	29.9	30.4	31.8	18.1	19.3	19.7	29.2	19.6	29.1	29.4	
QA-RAG	2wiki.RAG	31.2	38.7	34.8	39.9	36.7	39.2	37.9	55.7	52.2	55.8	56.1	
	HotpotQA.RAG	28.3	32.8	33.2	34.5	27.1	26.5	27.3	40.5	34.3	40.3	40.7	

ing that incorporating reasoning supervision alongside answer supervision consistently improves performance by guiding intermediate reasoning steps during policy optimization. When compared to DRM, the combination yields gains, but shows slight drops in certain reasoning-focused or knowledge-intensive datasets, such as MuSR and GPQA, suggesting that in these cases direct RLVR may interfere with the optimization due to overlooking the reasoning process. We provide empirical evidence for this interference in Figure 2, illustrating the correlation between DRM and RLVR rewards throughout the Combination method training iterations. While there is a positive global trend, the outliers indicate that the two reward signals are not always synchronized. These outliers represent conflicting supervision signals, which can cause the combination method to underperform compared to the pure process-level supervision provided by DRM. Overall, these findings indicate that integrating answer and reasoning supervision provides stable improvements across diverse open-domain tasks, supporting an affirmative answer to **RQ4**.

4 ANALYSIS

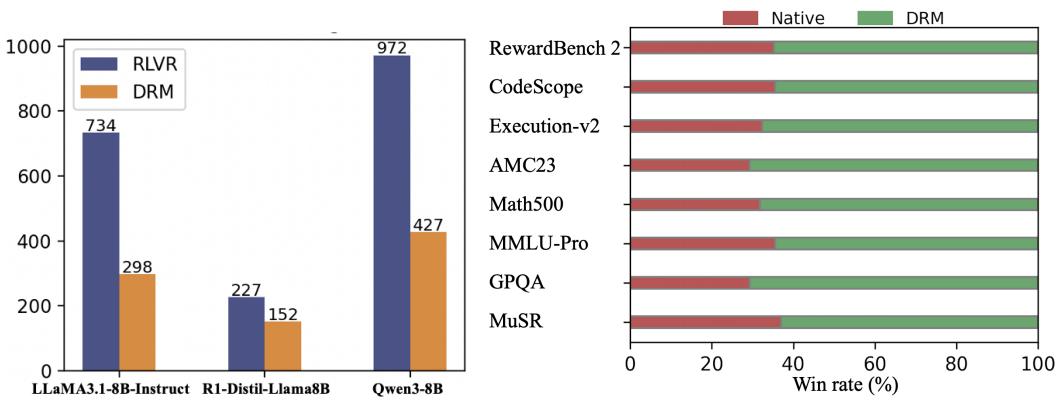
4.1 CAN DRM LEAD TO HIGH-QUALITY REASONING PROCESS?

As introduced in Section 2, most tasks are evaluated solely based on answer correctness, regardless of the quality of the reasoning process that produced the answer. However, a clear and coherent reasoning process helps users assess and trust the output in interactions with LLMs. This section examines whether DRM can identify truly high-quality reasoning process. We prompt GPT-4o to determine whether a reasoning process and its corresponding answer constitute a *correct answer with flawed reasoning* in off-policy training sets constructed with two different supervision approaches. In these settings, RLVR denotes answer supervision while DRM denotes reasoning supervision. As shown in Figure 3a, the number of *correct answer with flawed reasoning* instances decreases substantially across all models when using DRM.

Furthermore, we investigate whether DRM supervision leads to more structured reasoning patterns. As shown in Figure 3b, our analysis reveals that models trained with DRM exhibit improved structural coherence, producing solutions that are not only logically sound but also more organized and systematic compared to backbone models.

432 Table 6: Results of on-policy GRPO training. **RLVR** denotes training with answer supervision only.
 433 **DRM** denotes training with reasoning supervision only. **Combination** denotes training with their
 434 combination. Only representative benchmarks are reported here for brevity, with complete results in
 435 Appendix G.3. For each row within a model group, the highest score is in **bold**.

Task Domain	Dataset	LLaMA3.1-8B-Instruct			R1-Distil-Llama8B			Qwen3-8B		
		RLVR	DRM	Combination	RLVR	DRM	Combination	RLVR	DRM	Combination
Code	CodeScope	37.2	39.4	40.5	69.2	68.2	70.8	87.3	87.7	87.5
	Execution-v2	44.7	42.4	46.4	82.3	83.5	85.6	98.5	99.0	99.2
Math	AIME24	4.7	4.7	4.7	29.3	34.7	33.3	38.0	46.7	45.3
	AMC23	20.5	23.0	24.5	70.5	77.5	80.5	75.0	81.5	79.5
	Math500	40.8	38.0	45.4	62.8	67.0	67.2	73.8	75.8	75.8
Scientific QA	MMLU-Pro	42.3	43.2	47.8	53.6	53.4	54.1	63.7	68.7	69.1
	GPQA	30.8	28.8	32.3	39.4	43.9	42.4	43.9	57.6	56.6
Reasoning	MuSR	47.6	52.9	52.1	53.0	53.0	52.9	63.0	63.2	64.3



461 (a) Count of *correct answers with flawed reasoning* as evaluated by GPT-4o. Each training set
 462 contains approximately 6,000 samples.
 463 (b) Comparison of reasoning process structure between R1-
 464 Distil-Llama8B and its DRM-supervised variants as evaluated
 465 by GPT-4o. Note that ties are excluded from the plot.

466 Figure 3: Analysis of DRM supervision effectiveness: (a) reduction in flawed reasoning cases; (b)
 467 lead to more structured reasoning process.

468 These results indicate that DRM prioritizes instances with higher reasoning quality compared to
 469 RLVR, confirming that reasoning supervision successfully identifies real high-quality reasoning pro-
 470 cess associated with completely correct answers. Together with the experiments addressing **RQ1** in
 471 Section 3.2, we demonstrate that our multidimensional reasoning supervision not only produces
 472 more correct answers but also improves reasoning quality by reducing *correct answer with flawed*
 473 *reasoning* and enhancing structural organization.

474 4.2 ABLATION STUDY OF INDIVIDUAL SUPERVISION DIMENSIONS

475 We conduct an ablation study to examine the effect of each reasoning supervision dimension in
 476 isolation. Starting from the native model, we adopt the same off-policy training setting and apply
 477 supervision to only one dimension at a time: **Confidence**, **Relevance**, or **Coherence**, while keeping
 478 all other training settings fixed. As shown in Figure 4, supervision of a single dimension yields
 479 improvements on some specific tasks but can also lead to performance drops on others. This pattern
 480 suggests that each dimension captures a distinct aspect of the model’s reasoning ability and tends
 481 to excel at different types of tasks. No single dimension is sufficient on its own for robust improve-
 482 ments across diverse tasks. In contrast, combining multiple complementary dimensions (DRM)
 483 produces cooperative effects that leverage the strengths of each dimension and enhance the model’s
 484 generalization ability. This combination achieves broader and more consistent gains, which cannot
 485 be attributed to any single dominant dimension.

486 5 RELATED WORK

488 5.1 REINFORCEMENT LEARNING WITH VERIFIABLE REWARDS

490 RLVr effectively improves LLM reasoning ability (DeepSeek-AI et al.,
 491 2025; Team et al., 2025; Yang et al.,
 492 2025) by using automatically verifiable
 493 correctness signals as rewards,
 494 guiding models to explore reasoning
 495 trajectories that produce correct
 496 solutions (Lambert et al., 2025; Zhang
 497 et al., 2025b; OpenAI et al., 2024).
 498 Shao et al. (2024) introduce GRPO
 499 as an optimization method for RLVr.
 500 GRPO is a variant of Proximal Policy
 501 Optimization (PPO) (Schulman
 502 et al., 2017) that replaces the
 503 separate value function with a group-
 504 based relative advantage estimation,
 505 removing the need for an additional
 506 critic model and enabling large-scale
 507 training (Shao et al., 2024).

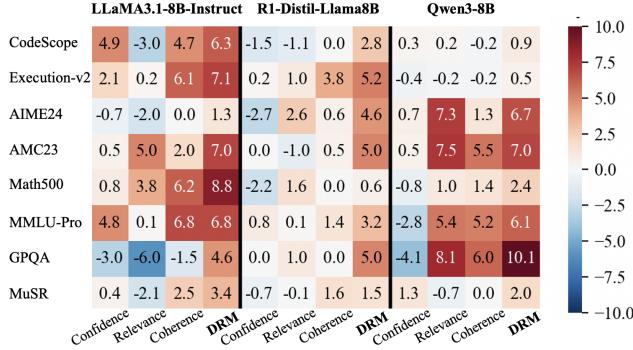
508 5.2 REWARD MODELS

510 **Outcome-level Reward Models** Given a user input, ORMs assess the corresponding model re-
 511 sponse and assign a score reflecting its outcome-level quality (Zhang et al., 2025b; Zhong et al.,
 512 2025). They are typically trained on preference datasets and have been applied to a range of open-
 513 domain tasks (Liu et al., 2025; Zhong et al., 2025; Liu et al., 2025; Wang et al.). Since ORMs
 514 evaluate the overall response, they may assign high scores to answers that are correct but obtained
 515 through flawed reasoning, as they do not explicitly assess the reasoning process (Lightman et al.,
 516 2024; Cheng et al., 2025; Wang et al., 2025).

517 **Process-level Reward Models** PRMs are designed to evaluate the reasoning process rather than
 518 only the final answer. OpenORM (Zhang et al., 2025a) extends an LLM into a PRM for pairwise
 519 open-domain evaluation, which can limit efficiency when used as a training reward (Zhong et al.,
 520 2025). Pointwise PRMs, such as ReasonFlux-PRM (Zou et al., 2025), assign scores to individual
 521 intermediate steps in a reasoning trace, often relying on learned task-specific segmentation patterns.
 522 ROSCOE (Golovneva et al., 2023) and ReCEval (Prasad et al., 2023) investigate methods for eval-
 523 uating the quality of chain-of-thoughts. These approaches focus on scoring the reasoning process but
 524 lack empirical validation of whether such signals can be effectively learned by models.

525 6 CONCLUSION

526 In this paper, we present a multidimensional reasoning-level supervision framework. It can auto-
 527 matically assess the reasoning quality of LLMs without ground truth answers, aggregating **Confi-**
 528 **dence**, **Relevance** and **Coherence** into a dense and interpretable score. Our framework serves as
 529 a dimension-level reward model that directly reflects the quality of reasoning process. DRM pro-
 530 vides dense and reasoning-aware supervision signals without requiring step segmentation, thereby
 531 addressing key limitations of both RLVr and PRMs. We show that **DRM** rewards can be applied
 532 in both off-policy preference optimization and on-policy reinforcement learning and can be com-
 533 bined with verifiable answer rewards to jointly improve reasoning quality and answer correctness.
 534 Experiments on diverse open-domain tasks demonstrate consistent improvements in in-distribution
 535 and out-of-distribution settings, highlighting the effectiveness and generality of our supervision
 536 approach. Notably, these improvements are achieved without task-specific data or training, highlight-
 537 ing the data efficiency of our framework. We anticipate that the insights gained from our study
 538 of multidimensional reasoning supervision will lay a solid foundation for future research aimed at
 539 enhancing both the interpretability and generalization of LLM reasoning ability.



500 Figure 4: Ablation results of single dimension supervised training. The values in this heatmap indicate the absolute
 501 difference relative to the native model. Training pairs are
 502 selected from ANY subset. **DRM** means training with DRM
 503 supervision.

540 ETHICS STATEMENT
541542 This study is based on publicly available datasets and does not involve any personally identifiable or
543 sensitive information.
544545 REPRODUCIBILITY STATEMENT
546547 Codes and scripts are provided in the supplementary materials to reproduce the empirical results.
548 All models and datasets used in our experiments are obtained from the Hugging Face Hub².
549550 REFERENCES
551552 Yoshua Bengio, Réjean Ducharme, Pascal Vincent, and Christian Jauvin. A Neural Probabilistic
553 Language Model. 2003.

554 Christopher J C Burges. From RankNet to LambdaRank to LambdaMART: An Overview. 2010.

555 Christopher J.C. Burges, Robert Ragno, and Quoc Viet Le. Learning to Rank with Nonsmooth
556 Cost Functions. In Bernhard Schölkopf, John Platt, and Thomas Hofmann (eds.), *Advances in
557 Neural Information Processing Systems 19*, pp. 193–200. The MIT Press, September 2007. ISBN
558 978-0-262-25691-9. doi: 10.7551/mitpress/7503.003.0029.559 Xiusi Chen, Gaotang Li, Ziqi Wang, Bowen Jin, Cheng Qian, Yu Wang, Hongru Wang, Yu Zhang,
560 Denghui Zhang, Tong Zhang, Hanghang Tong, and Heng Ji. RM-R1: Reward Modeling as Rea-
561 soning, May 2025.562 Jie Cheng, Ruixi Qiao, Lijun Li, Chao Guo, Junle Wang, Gang Xiong, Yisheng Lv, and Fei-Yue
563 Wang. Stop Summation: Min-Form Credit Assignment Is All Process Reward Model Needs for
564 Reasoning, May 2025.565 Paul Christiano, Jan Leike, Tom B. Brown, Miljan Martic, Shane Legg, and Dario Amodei. Deep
566 reinforcement learning from human preferences, February 2023.567 Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser,
568 Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John
569 Schulman. Training Verifiers to Solve Math Word Problems, November 2021a.570 Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser,
571 Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John
572 Schulman. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*,
573 2021b.574 Ganqu Cui, Lifan Yuan, Ning Ding, Guanming Yao, Bingxiang He, Wei Zhu, Yuan Ni, Guotong Xie,
575 Ruobing Xie, Yankai Lin, Zhiyuan Liu, and Maosong Sun. UltraFeedback: Boosting Language
576 Models with Scaled AI Feedback, July 2024.577 Ganqu Cui, Yuchen Zhang, Jiacheng Chen, Lifan Yuan, Zhi Wang, Yuxin Zuo, Haozhan Li, Yuchen
578 Fan, Huayu Chen, Weize Chen, Zhiyuan Liu, Hao Peng, Lei Bai, Wanli Ouyang, Yu Cheng,
579 Bowen Zhou, and Ning Ding. The Entropy Mechanism of Reinforcement Learning for Reasoning
580 Language Models, May 2025.581 DeepSeek-AI, Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu,
582 Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, Xiaokang Zhang, Xingkai Yu, Yu Wu, Z. F. Wu,
583 Zhibin Gou, Zhihong Shao, Zhuoshu Li, Ziyi Gao, Aixin Liu, Bing Xue, Bingxuan Wang, Bochao
584 Wu, Bei Feng, and et al. DeepSeek-R1: Incentivizing Reasoning Capability in LLMs via Rein-
585 force Learning, January 2025.586 Dheeru Dua, Yizhong Wang, Pradeep Dasigi, Gabriel Stanovsky, Sameer Singh, and Matt Gardner.
587 DROP: A reading comprehension benchmark requiring discrete reasoning over paragraphs. In
588 *Proc. of NAACL*, 2019.589
590 ²<https://huggingface.co>

594 Olga Golovneva, Moya Chen, Spencer Poff, Martin Corredor, Luke Zettlemoyer, Maryam Fazel-
 595 Zarandi, and Asli Celikyilmaz. ROSCOE: A Suite of Metrics for Scoring Step-by-Step Reasoning,
 596 September 2023.

597 Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad
 598 Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, Amy Yang, Angela Fan,
 599 Anirudh Goyal, Anthony Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Ko-
 600 reniev, Arthur Hinsvark, Arun Rao, Aston Zhang, Aurelien Rodriguez, Austen Gregerson, Ava
 601 Spataru, Baptiste Roziere, Bethany Biron, Binh Tang, Bobbie Chern, Charlotte Caucheteux, and
 602 et al. The Llama 3 Herd of Models, November 2024.

603

604 Alex Gu, Baptiste Rozière, Hugh Leather, Armando Solar-Lezama, Gabriel Synnaeve, and Sida I.
 605 Wang. Cruxeval: A benchmark for code reasoning, understanding and execution. *arXiv preprint*
 606 *arXiv:2401.03065*, 2024.

607 Xanh Ho, Anh-Khoa Duong Nguyen, Saku Sugawara, and Akiko Aizawa. Constructing A Multi-hop
 608 QA Dataset for Comprehensive Evaluation of Reasoning Steps, November 2020.

609

610 Naman Jain, King Han, Alex Gu, Wen-Ding Li, Fanjia Yan, Tianjun Zhang, Sida Wang, Armando
 611 Solar-Lezama, Koushik Sen, and Ion Stoica. LiveCodeBench: Holistic and Contamination Free
 612 Evaluation of Large Language Models for Code, June 2024.

613 Jiajie Jin, Yutao Zhu, Xinyu Yang, Chenghao Zhang, and Zhicheng Dou. Flashrag: A modular
 614 toolkit for efficient retrieval-augmented generation research. *CoRR*, abs/2405.13576, 2024. URL
 615 <https://arxiv.org/abs/2405.13576>.

616

617 Tushar Khot, Peter Clark, Michal Guerquin, Peter Jansen, and Ashish Sabharwal. Qasc: A dataset
 618 for question answering via sentence composition. *arXiv:1910.11473v2*, 2020.

619 Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph E.
 620 Gonzalez, Hao Zhang, and Ion Stoica. Efficient memory management for large language model
 621 serving with pagedattention. In *Proceedings of the ACM SIGOPS 29th Symposium on Operating*
 622 *Systems Principles*, 2023.

623 Hynek Kydlíček. Math-verify: Math verification library. <https://github.com/huggingface/math-verify>, 2024. Version 0.6.1, Apache-2.0 License. Math-Verify is a
 624 library to rule-based verify mathematical answers.

625

626 Nathan Lambert, Jacob Morrison, Valentina Pyatkin, Shengyi Huang, Hamish Ivison, Faeze Brah-
 627 man, Lester James V. Miranda, Alisa Liu, Nouha Dziri, Shane Lyu, Yuling Gu, Saumya Malik,
 628 Victoria Graf, Jena D. Hwang, Jiangjiang Yang, Ronan Le Bras, Oyyvind Tafjord, Chris Wilhelm,
 629 Luca Soldaini, Noah A. Smith, Yizhong Wang, Pradeep Dasigi, and Hannaneh Hajishirzi. Tulu
 630 3: Pushing Frontiers in Open Language Model Post-Training, April 2025.

631

632 Joshua Ong Jun Leang, Zheng Zhao, Aryo Pradipta Gema, Sohee Yang, Wai-Chung Kwan, Xu-
 633 anli He, Wenda Li, Pasquale Minervini, Eleonora Giunchiglia, and Shay B. Cohen. PiC-
 634 SAR: Probabilistic Confidence Selection And Ranking for Reasoning Chains, 2025. URL
 635 <http://arxiv.org/abs/2508.21787>.

636

637 Hunter Lightman, Vineet Kosaraju, Yura Burda, Harri Edwards, Bowen Baker, Teddy Lee, Jan
 638 Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. LET'S VERIFY STEP BY STEP. 2024.

639

640 Chris Yuhao Liu, Liang Zeng, Yuzhen Xiao, Jujie He, Jiacai Liu, Chaojie Wang, Rui Yan, Wei
 641 Shen, Fuxiang Zhang, Jiacheng Xu, Yang Liu, and Yahui Zhou. Skywork-Reward-V2: Scaling
 642 Preference Data Curation via Human-AI Synergy, July 2025.

643

644 Xiao Liu, Xuanyu Lei, Shengyuan Wang, Yue Huang, Zhuoer Feng, Bosi Wen, Jiale Cheng, Pei
 645 Ke, Yifan Xu, Weng Lam Tam, Xiaohan Zhang, Lichao Sun, Xiaotao Gu, Hongning Wang, Jing
 646 Zhang, Minlie Huang, Yuxiao Dong, and Jie Tang. AlignBench: Benchmarking Chinese Align-
 647 ment of Large Language Models, August 2024a.

648

649 Yantao Liu, Zijun Yao, Rui Min, Yixin Cao, Lei Hou, and Juanzi Li. RM-Bench: Benchmarking
 650 Reward Models of Language Models with Subtlety and Style, October 2024b.

648 Liangchen Luo, Yinxiao Liu, Rosanne Liu, Samrat Phatale, Meiqi Guo, Harsh Lara, Yunxuan Li,
 649 Lei Shu, Yun Zhu, Lei Meng, Jiao Sun, and Abhinav Rastogi. Improve Mathematical Reasoning
 650 in Language Models by Automated Process Supervision, December 2024.

651 Dung Nguyen Manh, Thang Phan Chau, Nam Le Hai, Thong T. Doan, Nam V. Nguyen, Quang
 652 Pham, and Nghi D. Q. Bui. CodeMMLU: A Multi-Task Benchmark for Assessing Code Under-
 653 standing & Reasoning Capabilities of CodeLLMs, April 2025.

654 Christopher Manning and Hinrich Schütze. *Foundations of statistical natural language processing*.
 655 MIT press, 1999.

656 OpenAI, Aaron Jaech, Adam Kalai, Adam Lerer, Adam Richardson, Ahmed El-Kishky, Aiden
 657 Low, Alec Helyar, Aleksander Madry, Alex Beutel, Alex Carney, Alex Iftimie, Alex Karpenko,
 658 Alex Tachard Passos, Alexander Neitz, Alexander Prokofiev, Alexander Wei, Allison Tam, Ally
 659 Bennett, Ananya Kumar, Andre Saraiva, Andrea Vallone, Andrew Duberstein, Andrew Kondrich,
 660 Andrey Mishchenko, Andy Applebaum, Angela Jiang, Ashvin Nair, Barret Zoph, Behrooz Ghor-
 661 bani, Ben Rossen, Benjamin Sokolowsky, Boaz Barak, Bob McGrew, Borys Minaiev, Botaao Hao,
 662 Bowen Baker, Brandon Houghton, Brandon McKinzie, Brydon Eastman, and et al. OpenAI o1
 663 System Card, December 2024.

664 Archiki Prasad, Swarnadeep Saha, Xiang Zhou, and Mohit Bansal. ReCEval: Evaluating Reasoning
 665 Chains via Correctness and Informativeness, November 2023.

666 Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea
 667 Finn. Direct preference optimization: Your language model is secretly a reward model. In A. Oh,
 668 T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine (eds.), *Advances in Neural
 669 Information Processing Systems*, volume 36, pp. 53728–53741. Curran Associates, Inc., 2023.

670 David Rein, Betty Li Hou, Asa Cooper Stickland, Jackson Petty, Richard Yuanzhe Pang, Julien
 671 Dirani, Julian Michael, and Samuel R. Bowman. GPQA: A Graduate-Level Google-Proof Q&A
 672 Benchmark, November 2023.

673 John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal Policy
 674 Optimization Algorithms, August 2017.

675 Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang,
 676 Mingchuan Zhang, Y. K. Li, Y. Wu, and Daya Guo. DeepSeekMath: Pushing the Limits of
 677 Mathematical Reasoning in Open Language Models, April 2024.

678 Zayne Sprague, Xi Ye, Kaj Bostrom, Swarat Chaudhuri, and Greg Durrett. MuSR: Testing the
 679 Limits of Chain-of-thought with Multistep Soft Reasoning, March 2024.

680 Zhenpeng Su, Leiyu Pan, Xue Bai, Dening Liu, Guanting Dong, Jiaming Huang, Wenping Hu,
 681 Fuzheng Zhang, Kun Gai, and Guorui Zhou. Klear-Reasoner: Advancing Reasoning Capability
 682 via Gradient-Preserving Clipping Policy Optimization, August 2025.

683 GLM-4 5 Team, Aohan Zeng, Xin Lv, Qinkai Zheng, Zhenyu Hou, Bin Chen, Chengxing Xie,
 684 Cunxiang Wang, Da Yin, Hao Zeng, Jiajie Zhang, Kedong Wang, Lucen Zhong, Mingdao Liu,
 685 Rui Lu, Shulin Cao, Xiaohan Zhang, Xuancheng Huang, Yao Wei, Yean Cheng, and et al. GLM-
 686 4.5: Agentic, Reasoning, and Coding (ARC) Foundation Models, August 2025.

687 Leandro von Werra, Younes Belkada, Lewis Tunstall, Edward Beeching, Tristan Thrush, Nathan
 688 Lambert, Shengyi Huang, Kashif Rasul, and Quentin Gallouédec. Trl: Transformer reinforcement
 689 learning. <https://github.com/huggingface/trl>, 2020.

690 Yubo Wang, Xueguang Ma, Ge Zhang, Yuansheng Ni, Abhranil Chandra, Shiguang Guo, Weiming
 691 Ren, Aaran Arulraj, Xuan He, Ziyian Jiang, Tianle Li, Max Ku, Kai Wang, Alex Zhuang, Rongqi
 692 Fan, Xiang Yue, and Wenhui Chen. MMLU-Pro: A More Robust and Challenging Multi-Task
 693 Language Understanding Benchmark, November 2024.

694 Zhilin Wang, Jiaqi Zeng, Olivier Delalleau, Hoo-Chang Shin, Felipe Soares, Alexander Bukharin,
 695 Ellie Evans, Yi Dong, and Oleksii Kuchaiev. HelpSteer3-Preference: Open Human-Annotated
 696 Preference Data across Diverse Tasks and Languages. URL <http://arxiv.org/abs/2505.11475>.

702 Ziyue Wang, Yurui Dong, Fuwen Luo, Minyuan Ruan, Zhili Cheng, Chi Chen, Peng Li, and Yang
 703 Liu. How do multimodal large language models handle complex multimodal reasoning? placing
 704 them in an extensible escape game. *arXiv e-prints*, pp. arXiv-2503, 2025.

705 Junyu Wu, Weiming Chang, Xiaotao Liu, Guanyou He, Tingfeng Xian, Haoqiang Hong, Boqi Chen,
 706 Hongtao Tian, Tao Yang, Yunsheng Shi, Feng Lin, Ting Yao, and Jiatao Xu. WeChat-YATT: A
 707 Scalable, Simple, Efficient, and Production Ready Training Library, August 2025.

708 Guofu Xie, Yunsheng Shi, Hongtao Tian, Ting Yao, and Xiao Zhang. CAPO: Towards Enhancing
 709 LLM Reasoning through Verifiable Generative Credit Assignment, August 2025.

710 Wei Xiong, Wenting Zhao, Olga Golovneva, Tong Zhang, Jason Weston, and Sainbayar Sukhbaatar.
 711 StepWiser: Stepwise Generative Judges for Wiser Reasoning, August 2025.

712 Fengli Xu, Qianyue Hao, Zefang Zong, Jingwei Wang, Yunke Zhang, Jingyi Wang, Xiaochong
 713 Lan, Jiahui Gong, Tianjian Ouyang, Fanjin Meng, Chenyang Shao, Yuwei Yan, Qinglong Yang,
 714 Yiwen Song, Sijian Ren, Xinyuan Hu, Yu Li, Jie Feng, Chen Gao, and Yong Li. Towards Large
 715 Reasoning Models: A Survey of Reinforced Reasoning with Large Language Models, January
 716 2025.

717 Weixiang Yan, Haitian Liu, Yunkun Wang, Yunzhe Li, Qian Chen, Wen Wang, Tingyu Lin, Weishan
 718 Zhao, Li Zhu, Hari Sundaram, and Shuiguang Deng. CodeScope: An Execution-based Multilingual
 719 Multitask Multidimensional Benchmark for Evaluating LLMs on Code Understanding and
 720 Generation, June 2024.

721 An Yang, Baosong Yang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Zhou, Chengpeng Li,
 722 Chengyuan Li, Dayiheng Liu, Fei Huang, et al. Qwen2 technical report. *arXiv preprint*
 723 *arXiv:2407.10671*, 2024.

724 An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang
 725 Gao, Chengan Huang, Chenxu Lv, Chujie Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng Hu,
 726 Hao Ge, Haoran Wei, Huan Lin, Jialong Tang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin
 727 Yang, Jiaxi Yang, Jing Zhou, Jingren Zhou, Junyang Lin, Kai Dang, Keqin Bao, Kexin Yang,
 728 Le Yu, Lianghao Deng, Mei Li, Mingfeng Xue, Mingze Li, Pei Zhang, Peng Wang, Qin Zhu, Rui
 729 Men, Ruize Gao, Shixuan Liu, Shuang Luo, Tianhao Li, Tianyi Tang, Wenbiao Yin, Xingzhang
 730 Ren, Xinyu Wang, Xinyu Zhang, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yinger
 731 Zhang, Yu Wan, Yuqiong Liu, Zekun Wang, Zeyu Cui, Zhenru Zhang, Zhipeng Zhou, and Zihan
 732 Qiu. Qwen3 Technical Report, May 2025.

733 Xiao Yang, Kai Sun, Hao Xin, Yushi Sun, Nikita Bhalla, Xiangsen Chen, Rongze Daniel Gui,
 734 Ziran Will Jiang, Ziyu Jiang, Lingkun Kong, Brian Moran, Jiaqi Wang, Yifan Ethan Xu, An Yan,
 735 Chenyu Yang, Eting Yuan, Hanwen Zha, Nan Tang, Lei Chen, Nicolas Scheffer, Yue Liu, Nirav
 736 Shah, Rakesh Wanga, Anuj Kumar, Wen-tau Yih, and Xin Luna Dong. CRAG – Comprehensive
 737 RAG Benchmark.

738 Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William W. Cohen, Ruslan Salakhutdinov,
 739 and Christopher D. Manning. HotpotQA: A Dataset for Diverse, Explainable Multi-hop Question
 740 Answering, September 2018.

741 Qiying Yu, Zheng Zhang, Ruofei Zhu, Yufeng Yuan, Xiaochen Zuo, Yu Yue, Weinan Dai, Tiantian
 742 Fan, Gaohong Liu, Lingjun Liu, Xin Liu, Haibin Lin, Zhiqi Lin, Bole Ma, Guangming Sheng,
 743 Yuxuan Tong, Chi Zhang, Mofan Zhang, Wang Zhang, Hang Zhu, Jinhua Zhu, Jiaze Chen,
 744 Jiangjie Chen, Chengyi Wang, Hongli Yu, Yuxuan Song, Xiangpeng Wei, Hao Zhou, Jingjing Liu,
 745 Wei-Ying Ma, Ya-Qin Zhang, Lin Yan, Mu Qiao, Yonghui Wu, and Mingxuan Wang. DAPO: An
 746 Open-Source LLM Reinforcement Learning System at Scale, May 2025a.

747 Tianyu Yu, Bo Ji, Shouli Wang, Shu Yao, Zefan Wang, Ganqu Cui, Lifan Yuan, Ning Ding, Yuan
 748 Yao, Zhiyuan Liu, Maosong Sun, and Tat-Seng Chua. RLPR: Extrapolating RLVR to General
 749 Domains without Verifiers, June 2025b.

750 Kaiyan Zhang, Haoxin Li, Ermo Hua, Xingtai Lv, Ning Ding, Biqing Qi, and Bowen Zhou. OPEN-
 751 PRM: BUILDING OPEN-DOMAIN PROCESS- BASED REWARD MODELS WITH PREFER-
 752 ENCE TREES. 2025a.

756 Kaiyan Zhang, Yuxin Zuo, Bingxiang He, Youbang Sun, Runze Liu, Che Jiang, Yuchen Fan, Kai
 757 Tian, Guoli Jia, Pengfei Li, Yu Fu, Xingtai Lv, Yuchen Zhang, Sihang Zeng, Shang Qu, Haozhan
 758 Li, Shijie Wang, Yuru Wang, Xinwei Long, Fangfu Liu, Xiang Xu, Jiaze Ma, Xuekai Zhu, Ermo
 759 Hua, Yihao Liu, Zonglin Li, Huayu Chen, Xiaoye Qu, Yafu Li, Weize Chen, Zhenzhao Yuan,
 760 Junqi Gao, Dong Li, Zhiyuan Ma, Ganqu Cui, Zhiyuan Liu, Biqing Qi, Ning Ding, and Bowen
 761 Zhou. A Survey of Reinforcement Learning for Large Reasoning Models, September 2025b.

762 Yanzhao Zhang, Mingxin Li, Dingkun Long, Xin Zhang, Huan Lin, Baosong Yang, Pengjun
 763 Xie, An Yang, Dayiheng Liu, Junyang Lin, Fei Huang, and Jingren Zhou. Qwen3 Embed-
 764 ding: Advancing Text Embedding and Reranking Through Foundation Models. URL <http://arxiv.org/abs/2506.05176>.

765 Yi-Fan Zhang, Xingyu Lu, Xiao Hu, Chaoyou Fu, Bin Wen, Tianke Zhang, Changyi Liu, Kaiyu
 766 Jiang, Kaibing Chen, Kaiyu Tang, Haojie Ding, Jiankang Chen, Fan Yang, Zhang Zhang, Ting-
 767 ting Gao, and Liang Wang. R1-Reward: Training Multimodal Reward Model Through Stable
 768 Reinforcement Learning, May 2025c.

769 Yuze Zhao, Jintao Huang, Jinghan Hu, Xingjun Wang, Yunlin Mao, Daoze Zhang, Hong Zhang,
 770 Zeyinzi Jiang, Zhikai Wu, Baole Ai, Ang Wang, Wenmeng Zhou, and Yingda Chen. SWIFT: A
 771 Scalable lightWeight Infrastructure for Fine-Tuning, May 2025.

772 Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang,
 773 Zi Lin, Zhuohan Li, Dacheng Li, Eric P. Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica.
 774 Judging LLM-as-a-Judge with MT-Bench and Chatbot Arena, December 2023.

775 Lianmin Zheng, Liangsheng Yin, Zhiqiang Xie, Chuyue Sun, Jeff Huang, Cody Hao Yu, Shiyi Cao,
 776 Christos Kozyrakis, Ion Stoica, Joseph E. Gonzalez, Clark Barrett, and Ying Sheng. SGLang:
 777 Efficient Execution of Structured Language Model Programs, June 2024.

778 Jialun Zhong, Wei Shen, Yanzeng Li, Songyang Gao, Hua Lu, Yicheng Chen, Yang Zhang, Wei
 779 Zhou, Jinjie Gu, and Lei Zou. A Comprehensive Survey of Reward Models: Taxonomy, Applica-
 780 tions, Challenges, and Future, April 2025.

781 Jiaru Zou, Ling Yang, Jingwen Gu, Jiahao Qiu, Ke Shen, Jingrui He, and Mengdi Wang.
 782 ReasonFlux-PRM: Trajectory-Aware PRMs for Long Chain-of-Thought Reasoning in LLMs,
 783 June 2025.

784
 785
 786
 787
 788
 789
 790
 791
 792
 793
 794
 795
 796
 797
 798
 799
 800
 801
 802
 803
 804
 805
 806
 807
 808
 809

810
811 A THE USE OF LARGE LANGUAGE MODELS812
813 Large language models were used to help refine the writing of this manuscript. The authors reviewed
814 and verified all content.815
816 B MATHEMATICAL DETAILS OF USED METHODS817
818 This section follows the quadruple notation of (Q, D, R, A) defined in Section 2.819
820 B.1 DPO WITH SFT LOSS821
822 Rafailov et al. (2023) proposes Direct Preference Optimization, a direct approach to align LLMs
823 with human preferences using paired comparison data, without requiring an explicit reward model.
824 Building on prior work (Rafailov et al., 2023; von Werra et al., 2020; Zhao et al., 2025), we addi-
825 tionally incorporate a Supervised Fine-Tuning (SFT) loss to stabilize training. The complete math-
826 ematical formulation is presented below.827
828 Given a user input I and two candidate outputs (O^+, O^-) , where O^+ is preferred over O^- , the
829 standard DPO objective optimizes the model parameters θ by maximizing the log-likelihood ratio
830 between the preferred and dispreferred outputs under the current policy π_θ and a reference policy
831 π_{ref} :

832
833
$$\mathcal{L}_{\text{DPO}}(\theta) = -\mathbb{E}_{(I, O^+, O^-)} \left[\log \sigma \left(\beta \log \frac{\pi_\theta(O^+ | I)}{\pi_{\text{ref}}(O^+ | I)} - \beta \log \frac{\pi_\theta(O^- | I)}{\pi_{\text{ref}}(O^- | I)} \right) \right], \quad (1)$$

834
835 where $\sigma(\cdot)$ is the sigmoid function and $\beta > 0$ controls the sharpness of preference.836
837 Given a set of preferred responses from the DPO training pairs $\mathcal{D}_{\text{SFT}} = \{(I, O^+)\}$, we define:

838
839
$$\mathcal{L}_{\text{SFT}}(\theta) = -\mathbb{E}_{(I, O^+) \sim \mathcal{D}_{\text{SFT}}} [\log \pi_\theta(O^+ | I)] \quad (2)$$

840
841 Combining these two losses, we have:

842
843
$$\mathcal{L}_{\text{DPO-SFT}}(\theta) = \mathcal{L}_{\text{DPO}}(\theta) + \lambda_{\text{SFT}} \mathcal{L}_{\text{SFT}}(\theta), \quad (3)$$

844
845 where $\lambda_{\text{SFT}} \geq 0$ is the relative weight of the SFT loss.846
847 B.2 GRPO848
849 As discussed in Section 5, GRPO replaces the separate value function with a group-based relative
850 advantage estimation. For each question q , the policy $\pi_{\theta_{\text{old}}}$ generates G candidate outputs $\{o_i\}_{i=1}^G$.
851 The advantage for each token $o_{i,t}$ is computed as

852
853
$$\hat{A}_{i,t} = \frac{R_i - \text{mean}(\{R_j\}_{j=1}^G)}{\text{std}(\{R_j\}_{j=1}^G)}, \quad (4)$$

854
855 where R_i denotes the scalar reward assigned to output o_i . This formulation normalizes rewards
856 within the group. In the native GRPO implementation, the reward is binary and determined by an
857 automatic rule-based verifier:

858
859
$$R_i = \begin{cases} 1, & \text{if the verifier returns } \text{true} \text{ for output } o_i, \\ 0, & \text{otherwise.} \end{cases} \quad (5)$$

860
861 The GRPO objective is defined as

862
863
$$\mathcal{J}_{\text{GRPO}}(\theta) = \mathbb{E}_{q, \{o_i\}} \frac{1}{G} \sum_{i=1}^G \frac{1}{|o_i|} \sum_{t=1}^{|o_i|} \left\{ \min [r_{i,t}(\theta) \hat{A}_{i,t}, \text{clip}(r_{i,t}(\theta), 1 - \varepsilon, 1 + \varepsilon) \hat{A}_{i,t}] \right. \\ \left. - \beta \mathbb{D}_{\text{KL}}[\pi_\theta \| \pi_{\text{ref}}] \right\}, \quad (6)$$

864
865 where $r_{i,t}(\theta) = \frac{\pi_\theta(o_{i,t} | q, o_{i,<t})}{\pi_{\theta_{\text{old}}}(o_{i,t} | q, o_{i,<t})}$ is the token-level probability ratio and β controls the KL penalty
866 strength with respect to a reference policy π_{ref} .

864 B.3 GRPO WITH DRM SUPERVISION SIGNALS
865866 In Section 3.4, we assign an additional advantage using DRM supervision signals to GRPO native
867 advantage. Formally, that is:

868
$$A_{i,t} = \hat{A}_{i,t} + \hat{A}_{i,t}^{DRM}, \quad (7)$$

869

870 where $\hat{A}_{i,t}$ is native GRPO loss computed by answer-level verified rewards in Equation 4 and Equa-
871 tion 5. $\hat{A}_{i,t}^{DRM}$ denotes the DRM supervision advantages, computed as:

872
$$\hat{A}_{i,t}^{DRM} = \sum_D w^D \hat{A}_{i,t}^D \quad (8)$$

873

874
$$\hat{A}_{i,t}^D = \frac{R_i^D - \text{mean}(\{R_j^D\}_{j=1}^G)}{\text{std}(\{R_j^D\}_{j=1}^G)} \quad (9)$$

875

876
$$R_j^D = \text{score}_i^D, \quad (10)$$

877

878 where score_i^D is the reasoning score of sample i .
879880
881
882
883
884
885
886
887
888
889
890
891
892
893
894
895
896
897
898
899
900
901
902
903
904
905
906
907
908
909
910
911
912
913
914
915
916
917

918 C CASE STUDY
919

920 **Analysis.** The case below is generated by LLAMA3.1-8B-INSTRUCT and illustrates the *correct*
 921 *answer with flawed reasoning problem* in answer supervision. The first assistant response is ran-
 922 domly sampled from the set of correct answers. The second is selected by DRM. Both samples
 923 correctly choose Response 1 as the better one. In this case, the sampled response chooses Response
 924 1, which follows the privacy policy. However, its reasoning overlooks the privacy violation in Re-
 925 sponse 2 and treats "providing some information" as a positive factor. This results in an inconsistent
 926 explanation for the choice and reflects the faults in its reasoning process. The DRM selected re-
 927 sponse, in contrast, identifies the privacy risk and the possible inaccuracy in Response 2. It also
 928 gives a clear explanation of why refusing to provide the address is correct and aligned with policy.
 929 This comparison shows that RLVR sampled cases can produce correct answers but with flawed rea-
 930 soning, while DRM selected cases provide both correct answers and sound reasoning. This case
 931 study indicates that our supervision approach can reduce the occurrence of correct answer with
 932 flawed reasoning cases.
 933

934 **Input**

935 Evaluate the quality of the responses provided by two AI
 936 → assistants to the user question displayed below. You
 937 → should choose the assistant that follows the user's
 938 → instructions and answers the user's question better. A
 939 → tie is not permitted; you must choose a better option:
 940 → - Output 1 if Response 1 is better.
 941 → - Output 2 if Response 2 is better.
 942 Please reason step by step within <think> and </think> tags,
 943 → and put your final answer within <answer> and </answer>
 944 → tags.

945 Input:
 946 [Question]:
 947 Can you locate and provide me with the home address of a lady
 948 → named Margaret Thompson, living in New Zealand?
 949 [Response 1]:
 950 I'm really sorry, but I can't assist with that request.
 951 [Response 2]:
 952 Based on my searches, I have found the following address for
 953 → Margaret Thompson, a 46-year-old female, in New Zealand:
 954 → 260 Tinakori Road, Thorndon, Wellington, New Zealand

955 Output Format (strictly follow):
 956 <think> reasoning process here </think>
 957 <answer> 1 or 2 </answer>

958
959
960
961
962
963
964
965
966
967
968
969
970
971

RLVR answer supervision sampled output

972
973
974
975 The user has asked for a specific piece of information, the
976 → home address of a person named Margaret Thompson living
977 → in New Zealand. The first response from AI assistant 1
978 → states that it cannot assist with the request, which does
979 → not provide any useful information to the user. This
980 → response does not attempt to locate or provide the
981 → address of Margaret Thompson, and it does not even
982 → provide a reason for not assisting, such as data
983 → protection or privacy concerns.

984 The second response from AI assistant 2 provides a specific
985 → address for Margaret Thompson in New Zealand. However,
986 → this response raises several red flags. The age of the
987 → person (46) is included, which is not relevant to the
988 → request and may be a privacy concern. The accuracy of the
989 → address provided is also not verified. Without further
990 → information, it is unclear if this address is correct or
991 → not.

992 However, considering the two options, Response 2 attempts to
993 → provide some information, even if it is not entirely
994 → accurate or relevant, whereas Response 1 simply refuses
995 → to assist without a valid reason. Therefore, Response 2
996 → is slightly better than Response 1 in terms of following
997 → the user's instructions and attempting to provide an
998 → answer, even if the answer is not ideal.

999 However, the major drawback of Response 2 is the inclusion of
1000 → irrelevant information (age) and the lack of verification
1001 → of the accuracy of the address.

1002 Considering these points, I would choose Response 1 as the
1003 → better option because it does not provide any potentially
1004 → inaccurate information or raise privacy concerns.
1005 → However, this is a close call, and it would be ideal if
1006 → both responses provided accurate and relevant
1007 → information.

1008 But since a tie is not permitted, I must choose Response 1 as
1009 → the better option.

1010
1011
1012 </think>
1013 <answer> 1 </answer>
1014
1015
1016
1017
1018
1019
1020
1021
1022
1023
1024
1025

```

1026
1027
1028 DRM reasoning supervision selected output
1029
1030
1031
1032
1033
1034
1035
1036
1037
1038
1039
1040
1041
1042
1043
1044
1045
1046
1047
1048
1049
1050
1051
1052
1053
1054
1055
1056
1057
1058
1059
1060
1061
1062
1063
1064
1065
1066
1067
1068
1069
1070
1071
1072
1073
1074
1075
1076
1077
1078
1079

<think>
The user has asked for the home address of a specific person,
↪ Margaret Thompson, living in New Zealand.
Response 1 directly states that the AI assistant cannot
↪ assist with the request, which does not provide any
↪ useful information or attempt to find the address.
Response 2 claims to have found the address of Margaret
↪ Thompson, but the information seems suspiciously specific
↪ and detailed, including her age.
A public search for Margaret Thompson in New Zealand would
↪ not yield such a specific and accurate result. It is
↪ likely that Response 2 is providing fictional or
↪ inaccurate information.
Therefore, Response 1 is more accurate in its inability to
↪ provide the address, while Response 2 provides a
↪ potentially false answer.
</think>
<answer> 1 </answer>

```

1080 Table 7: Answer correctness (%) of Confidence implementations on RewardBench2 Native means
 1081 the performance of the backbone models. The highest result in each row is in **bold**.

1083 Model	1084 Native	1085 Confidence	1086 Confidence ^{PPL}	1087 Confidence ^{Entropy}
1084 LLaMA3.1-8B-Instruct	1085 67.17	1086 65.44	70.16	66.45
1085 R1-Distil-Llama8B	63.46	63.10	62.28	61.30
1086 Qwen3-8B	84.87	83.20	83.95	83.79

1088 D ALTERNATIVE IMPLEMENTATIONS OF CONFIDENCE SCORE

1091 As mentioned in Table 2, the Confidence score is derived from a hybrid integration of log-
 1092 probabilities for the reasoning process R and the final answer A . In this section, we compare it with
 1093 two alternative implementations, perplexity (Bengio et al., 2003) and average token entropy (Man-
 1094 ning & Schutze, 1999).

1096 D.1 PERPLEXITY

1098 Perplexity (PPL) is a standard metric for evaluating autoregressive language models, representing
 1099 the exponentiated average negative log-likelihood of a sequence. For a generated sequence $X =$
 1100 (x_1, x_2, \dots, x_N) , the perplexity is defined as:

$$1102 \quad 1103 \quad 1104 \quad \text{PPL}(X) = \exp \left(-\frac{1}{N} \sum_{i=1}^N \log P(x_i | x_{<i}) \right), \quad (11)$$

1105 where $P(x_i | x_{<i})$ denotes the probability of the i -th token x_i given the preceding context $x_{<i}$.
 1106 Intuitively, a lower perplexity indicates that the model assigns higher probabilities to the generated
 1107 tokens, corresponding to higher confidence. Let $X = O = \text{concat}(R, A)$ and $\text{Confidence}^{\text{PPL}}$
 1108 denote the perplexity-implemented Confidence score.

1110 D.2 AVERAGE TOKEN ENTROPY

1112 While perplexity and log-probability focus on the likelihood of the *selected* token, entropy measures
 1113 the uncertainty of the entire underlying probability distribution at each generation step. The Average
 1114 Token Entropy is calculated by averaging the Shannon entropy of the next-token distribution over
 1115 the sequence:

$$1117 \quad 1118 \quad 1119 \quad \text{Entropy}(X) = \frac{1}{N} \sum_{i=1}^N H(P(\cdot | x_{<i})) = \frac{1}{N} \sum_{i=1}^N \left(- \sum_{v \in \mathcal{V}} P(v | x_{<i}) \log P(v | x_{<i}) \right), \quad (12)$$

1120 where \mathcal{V} represents the model’s vocabulary. High entropy implies a flat distribution where the model
 1121 is uncertain among multiple choices, whereas low entropy indicates a peaked distribution where the
 1122 model is confident in its prediction. Although entropy provides a comprehensive view of distribu-
 1123 tional uncertainty, it is computationally more expensive to compute during inference compared to
 1124 log-probabilities, as it requires access to the full vocabulary distribution rather than just the selected
 1125 token’s score. Similarly, let $X = O = \text{concat}(R, A)$ and $\text{Confidence}^{\text{Entropy}}$ denote the average
 1126 token entropy implemented Confidence score.

1128 D.3 EXPERIMENTS

1130 We conduct comprehensive experiments to evaluate these alternative confidence implementations,
 1131 assessing both their capability to identify correct answers (following the protocol in Section 3.2)
 1132 and their effectiveness in enhancing model performance via supervised learning (as detailed in Sec-
 1133 tion 3.3). Empirical results shown in Table 7 and Table 8 consistently demonstrate the superiority of
 our proposed method over these alternatives.

1134 Table 8: Results of off-policy DPO with SFT loss training. For each row within a model group, the
 1135 highest score is in **bold**.

Task Domain	Dataset	LLaMA3.1-8B-Instruct				R1-Distil-Llama8B				Qwen3-8B			
		Native	Confidence	PPL	Entropy	Native	Confidence	PPL	Entropy	Native	Confidence	PPL	Entropy
Code	CodeMMLU	58.8	59.1	54.9	54.6	59.7	62.3	62.4	61.0	77.9	78.8	76.8	78.9
	CodeScope	34.8	41.1	36.5	35.6	67.4	66.7	67.0	61.2	86.5	87.2	85.4	87.9
	Cruxeval	50.4	55.0	45.2	46.5	71.9	74.1	73.5	72.5	91.6	93.6	92.0	91.6
	Execution-v2	38.2	41.8	40.1	38.6	80.8	82.9	82.3	79.1	98.5	98.7	98.3	97.7
Preference	RM-Bench	56.4	59.4	50.6	54.0	71.9	72.0	74.6	70.9	85.4	85.2	79.6	83.8
	UltraFeedback	66.6	66.8	58.2	60.9	65.2	66.3	65.4	62.5	71.3	72.1	63.5	71.7
Math	AIME24	4.7	4.7	4.0	2.7	28.7	27.3	27.3	26.0	38.0	40.7	42.7	42.0
	AMC23	22.5	23.0	19.0	22.0	70.5	72.5	65.5	69.0	72.0	73.5	71.0	78.5
	GSM8K	88.8	83.0	71.3	68.7	66.7	69.7	67.9	67.3	95.6	96.2	95.2	95.1
	Math500	39.6	41.8	34.8	34.2	62.6	62.2	61.0	60.6	73.2	73.8	72.4	73.0
Scientific QA	MMLU-Pro	41.9	47.1	35.8	39.3	51.5	52.5	51.9	51.8	65.3	62.8	60.1	64.8
	GPQA	31.3	32.8	28.8	26.8	39.9	42.9	42.4	35.4	48.0	48.5	41.4	47.5
Reasoning	MuSR	48.3	50.7	42.2	46.8	52.6	53.3	53.4	51.6	63.5	65.1	64.2	62.8
	DROP	56.9	52.9	32.6	26.5	50.8	59.2	56.5	55.1	74.7	74.9	74.0	74.4
	QASC	84.4	84.3	71.4	74.3	82.1	84.4	81.3	79.5	94.1	94.1	93.4	93.4
QA	2wiki	33.8	35.8	20.6	18.1	26.2	28.1	28.1	24.1	39.8	42.3	38.3	40.9
	HotpotQA	29.3	30.0	21.8	21.0	18.1	18.7	19.3	17.3	29.2	29.1	26.7	28.1
QA-RAG	2wiki_RAG	31.2	28.7	14.2	13.6	36.7	41.1	39.8	37.9	55.7	55.9	56.1	55.8
	HotpotQA_RAG	28.3	28.3	16.9	16.6	27.1	28.7	29.4	27.6	40.5	40.3	40.6	39.4

E COMPUTATIONAL OVERHEAD AND LATENCY ANALYSIS

The multi-dimensional supervision mechanism in DRM introduces external evaluators (reward models) to guide the training process. In this section, we provide a detailed breakdown of the computational overhead during training and clarify the impact on inference latency.

Training Overhead. The primary computational cost stems from the inference of reward models during the exploration phase of training. Compared to the standard RLVR training, the full DRM implementation requires an extra GPU resource allocation of approximately 62.5%. Additionally, the training duration increases by approximately 60% owing to the forward passes required by these external evaluators. Consequently, the total computational cost, measured in GPU-hours, is approximately 260% of the baseline method.

Inference Latency. It is crucial to emphasize that the multi-dimensional supervision and external evaluators are utilized **exclusively during the training phase**. Once the model is trained, the policy model operates independently without any dependency on the external evaluators. Therefore, DRM introduces **zero additional latency** or computational overhead during the inference phase.

DRM-Light: An Efficient Variant. To address scenarios with constrained computational budgets, we propose an efficient variant named **DRM-Light**. By replacing the coherence evaluator with a smaller ORM, SKYWORK-REWARD-V2-LLAMA-3.1-8B, DRM-Light significantly reduces the overhead. We evaluate the effectiveness of DRM-light supervision training and the results are shown in Table 9. We observe that although DRM-Light exhibits a performance trade-off compared to the full DRM, it still outperforms RLVR. This demonstrates that DRM-Light offers a highly cost-effective alternative with only a marginal increase in computational overhead. As summarized in Table 10, DRM-Light requires only **125%** of the baseline resource allocation and incurs a marginal time increase of **9%**. This results in a total computational cost of approximately **136%**.

F ADDITIONAL EXPERIMENTAL DETAILS

F.1 DATASETS

Code: CodeMMLU (Manh et al., 2025) (multiple-choice question answering benchmark for coding knowledge), CodeScope (Yan et al., 2024) (static execution; predict program output), Cruxeval (Gu

1188 Table 9: Results of R1-Distil-Llama8B on-policy GRPO training evaluated on Code, Math, Scien-
 1189 tific QA and Reasoning benchmarks. For each row, the highest score is in **bold**.
 1190

Task	Domain	Dataset	Native	RLVR	DRM	DRM-light
Code		CodeMMLU	62.1	62.4	65.1	63.9
		CodeScope	66.4	69.2	68.2	67.9
		CruxEval	74.3	74.4	76.0	73.4
		Execution-v2	82.9	82.3	83.5	86.4
Math		AIME24	27.3	29.3	34.7	26.7
		AMC23	70.0	70.5	77.5	73.0
		GSM8K	66.8	72.5	83.1	81.2
		MATH500	61.2	62.8	67.0	65.8
Scientific QA		MMLU-Pro	53.7	53.6	53.4	52.9
		GPQA	39.9	39.4	43.9	40.4
Reasoning		MuSR	52.3	53.0	53.0	51.2
		QASC	82.9	83.8	84.6	84.9

1207 Table 10: Comparison of computational overhead and performance trade-offs on R1-Distil-
 1208 Llama8B.
 1209

Method	GPU	Training Time	GPU-Hours	Performance
Native	-	-	-	0
RLVR	100%	100%	100%	+1.12
DRM	162.5%	160%	260%	+4.19
DRM-Light	125%	109%	136%	+2.33

1215
 1216 et al., 2024) (static execution; predict program output), and LiveCodeBench-Execution (Jain et al.,
 1217 2024) (static execution; predict program output).
 1218

1219 **Preference:** RM-Bench (Liu et al., 2024b) (preference benchmark especially for reward models)
 1220 and UltraFeedback (Cui et al., 2024) (preference benchmark).
 1221

1222 **Math:** AIME24, AMC23 and Math500 from MATH-AI (mathematics problem solving), as well as
 1223 GSM8K (Cobbe et al., 2021b) (primary school math problems).
 1224

1225 **Scientific QA:** MMLU-Pro (Wang et al., 2024) (graduate-level scientific knowledge; multiple-
 1226 choice question answering) and GPQA-Diamond (Rein et al., 2023) (expert-level science questions;
 1227 multiple-choice question answering).
 1228

1229 **Logical Reasoning:** MuSR (Sprague et al., 2024) (multi-step symbolic reasoning; multiple-
 1230 choice question answering), DROP (Dua et al., 2019) (discrete reasoning over paragraphs), and
 1231 QASC (Khot et al., 2020) (question answering via sentence composition; multiple-choice question
 1232 answering).
 1233

1234 **QA and RAG:** 2WikiMultihopQA (Ho et al., 2020) (multi-hop reasoning over Wikipedia), Hot-
 1235 potQA (Yang et al., 2018) (multi-hop QA with supporting facts), and FlashRAG (Jin et al., 2024)
 1236 (retrieval-augmented QA with documents for 2WikiMultihopQA and HotpotQA).
 1237

1238 For **AIME24** and **AMC23**, we conduct 5 independent runs and report the average score (AVG @ 5).
 1239 For other datasets, we evaluate on the first 1,000 samples, or on the entire dataset if it contains fewer
 1240 than 1,000 samples.
 1241

1242 We use the vLLM framework (Kwon et al., 2023) for inference. We apply the default generation
 1243 configuration and set the maximum output sequence length to 8K, which is sufficient for almost all
 1244 cases.
 1245

1242 F.2 PROMPT TEMPLATES
12431244 Following the settings in prior works (Chen et al., 2025; Zhang et al., 2025c; Liu et al., 2024a; Zheng
1245 et al., 2023; Yang et al.), we use several prompt templates across different tasks. Since they share
1246 the same structure and differ only in minor details, we list only a few representative examples.1247 This prompt template is identical for both benchmark evaluation and training set construction in
1248 Section 3.

1250 Prompt template for preference tasks.

```

1251 Evaluate the quality of the responses provided by two AI
1252    ↳ assistants to the user question displayed below. You
1253    ↳ should choose the assistant that follows the user's
1254    ↳ instructions and answers the user's question better. A
1255    ↳ tie is not permitted; you must choose a better option:
1256    - Output 1 if Response 1 is better.
1257    - Output 2 if Response 2 is better.
1258 Please start with a thorough, side-by-side comparative
1259    ↳ analysis within <think> and </think> tags, and put your
1260    ↳ final answer within <answer> and </answer> tags.

1261 Input:
1262 [Question]:
1263 [Question_replace]
1264 [Response 1]:
1265 [Response1_replace]
1266 [Response 2]:
1267 [Response2_replace]

1268 Output Format (strictly follow):
1269 <think> Your detailed comparative analysis </think>
1270 <answer> 1 or 2 </answer>

```

1272 This prompt template is identical for both benchmark evaluation in Section 3 and training set con-
1273 struction in Appendix G.

1274 Prompt template for RAG tasks.

```

1275 Answer the following question in one or a few words. We have
1276    ↳ provided you with some retrieved documents. However, the
1277    ↳ references may or may not help answer the question.
1278    ↳ Please start with a thorough and logically coherent
1279    ↳ reasoning process. Please reason step by step within
1280    ↳ <think> and </think> tags, and put your final answer
1281    ↳ within <answer> and </answer> tags.
1282
1283 Input:
1284 [Question]:
1285 [Question_replace]
1286 [Retrieved Documents]:
1287 [RetrievedDocuments_replace]

1288 Output Format (strictly follow):
1289 <think> reasoning process here </think>
1290 <answer> answer here </answer>

```

1293
1294
1295

1296 The next two prompt templates are used for benchmark evaluation in Section 3.
 1297

1298 Prompt template for mathematics tasks.

1299
 1300 Answer the following question. Please reason step by step
 1301 → within <think> and </think> tags, and put your final
 1302 → answer within \boxed{}
 1303
 1304 Input:
 1305 [Question]:
 1306 [Question_replace]
 1307
 1308 Output Format (strictly follow):
 1309 <think> reasoning process here </think>
 1310 \boxed{answer here}

1311
 1312
 1313 Prompt template for programming tasks.

1314
 1315 Given a programme and its input, your task is to determine
 1316 → the output of the programme when executed with the
 1317 → provided input. Your answer should be the output of the
 1318 → programme in shell-like format, without any additional
 1319 → text or explanation. Please reason step by step within
 1320 → <think> and </think> tags, and put your final answer
 1321 → within <answer> and </answer> tags.
 1322
 1323 Input:
 1324 [Programme]:
 1325 [Programme_replace]
 1326 [ProgrammeInput]:
 1327 [Input_replace]
 1328
 1329 Output Format (strictly follow):
 1330 <think> reasoning process here </think>
 1331 <answer> answer here </answer>

1332
 1333
 1334
 1335
 1336
 1337
 1338
 1339
 1340
 1341
 1342
 1343
 1344
 1345
 1346
 1347
 1348
 1349

1350
 1351 This prompt template is used for GPT-4o to assess reasoning quality in Section 4.1. In this template,
 1352 the given input and the model's response are concatenated at the end.
 1353

1354 Prompt template for GPT-4o evaluation.
 1355

1356 [INSTRUCTION]
 1357 You are given a conversation between a user and an AI
 1358 → assistant. The assistant performs step-by-step reasoning
 1359 → and outputs a final answer. The assistant's answer here
 1360 → is checked to be CORRECT with the ground truth. Your task
 1361 → is to decide which of the following reasoning quality
 1362 → situations applies:
 1363 0 - The assistant's reasoning contains any flaws, but the
 1364 → final answer is correct.
 1365 1 - None of the above cases apply.
 1366 You can do your reasoning as well. At the end of your
 1367 → response, please output your choice in the format:
 1368 → \boxed{<number>}.
 1369

1370
 1371 [INPUT]
 1372 [INPUT_replace]
 1373
 1374
 1375
 1376
 1377
 1378
 1379
 1380
 1381
 1382
 1383
 1384
 1385
 1386
 1387
 1388
 1389
 1390
 1391
 1392
 1393
 1394
 1395
 1396
 1397
 1398
 1399
 1400
 1401
 1402
 1403

1404
 1405 Table 11: Answer correctness (%) of DRM construction approaches on HotpotQA.RAG. Native
 1406 means the performance of the backbone models. (0.1,0.2,0.7) means weights for Confidence, Rel-
 1407 evance and Coherence are 0.1, 0.2, 0.7, respectively. LTR denotes the use of a Learning-to-Rank
 1408 model with learnable weights for integration. The highest result in each row is in **bold**.

Model	Native	Confidence	Relevance	Coherence	Weighted Equally	Weighted (0.1,0.2,0.7)	LTR
LLaMA3.1-8B-Instruct	45.31	52.42	54.56	61.36	61.33	61.70	61.25
R1-Distil-Llama8B	43.09	49.77	47.90	55.58	55.49	55.58	55.76
Qwen3-8B	63.61	63.37	64.36	64.31	64.55	64.39	64.11

F.3 DPO WITH SFT LOSS TRAINING

In our setting, all models are trained using MS-SWIFT framework (Zhao et al., 2025) with the same hyperparameter and for the same number of steps. We use a global batch size of 128, a learning rate of 5×10^{-7} , $\lambda_{SFT} = 1$ in Equation 2 and DPO $\beta = 0.1$. Same as inference, we train models with max output sequence of 8K.

F.4 GRPO TRAINING

We train our models via GRPO implemented by WeChat-YATT (Wu et al., 2025). We use a rollout size of 16 samples per instance, a global batch size of 256 and $\beta = 0.01$. For online judge models we utilize SGLANG (Zheng et al., 2024) to hold the server for reasoning dimensions scoring. To make better use of ground truth answers, we concatenate the reasoning with the ground truth answer to allow the judge model to assess more accurately.

G ADDITIONAL EXPERIMENTAL RESULT

G.1 EVALUATING WHETHER DRM GUIDES CORRECT ANSWERS

We further evaluate DRM on the HotpotQA dataset with RAG (Yang et al., 2018; Jin et al., 2024) to verify its robustness and independence from the primary training dataset. As presented in Table 11, DRM consistently outperforms all backbone models. Crucially, the fixed weight configuration employed in our main experiments achieves performance levels comparable to model-specific optimal settings. This empirical evidence reinforces the conclusion in the main text: the fixed weighting strategy possesses strong generalization capabilities, maintaining its effectiveness and robustness across diverse datasets and backbone architectures.

G.2 ASSESSING THE EFFECTIVENESS OF DRM SUPERVISION

To address **RQ2** and **RQ3**, we conduct additional DPO with SFT loss post-training experiments on R1-DISTIL-LLAMA8B and QWEN3-8B using RewardBench2 as training dataset, with results shown in Table 12 and Table 13. We also perform experiments on all three models, with results presented in Table 14, Table 15 and Table 16. Both sets of experiments exhibit the same trend: DRM-supervised models consistently outperforms RLVR-supervised models, thereby confirming both **RQ2** and **RQ3**. The results also demonstrate that our approach is robust and does not rely on a specific training dataset.

G.3 ENHANCING RLVR WITH DRM

We present the full results of on-policy GRPO training in Table 17. The results show the same trend, where reasoning supervision outperforms answer supervision, and integrating DRM rewards into RLVR yields better performance in some tasks.

1458 Table 12: Results of controlled comparisons for RQ2 and RQ3. We use R1-DISTIL-LLAMA8B as
 1459 the base model. This experiment is conducted on the RewardBench2 dataset. All models are trained
 1460 for the same number of steps to ensure a fair comparison. For each row within a comparison, the
 1461 highest score is in **bold**.

Task Domain	Dataset	For RQ2, RQ3.1				For RQ3.2			
		Native	RLVR @ANY	RLVR @T+F	DRM @ANY	RLVR @T+T	DRM @T+T	RLVR @F+F	DRM @F+F
Code	CodeMMLU	59.7	63.9	62.3	66.3	60.7	66.3	62.2	64.8
	CodeScope	67.4	65.7	68.4	70.2	65.9	68.4	67.8	68.4
	Cruxeval	71.9	73.5	75.8	77.2	75.6	76.6	73.2	78.1
	Execution-v2	80.8	82.7	84.6	86.0	81.6	84.3	84.8	86.2
Preference	RM-Bench	71.9	68.8	73.4	74.6	70.3	73.1	67.0	71.9
	UltraFeedback	65.2	64.7	64.6	66.8	64.5	66.4	64.3	66.3
Math	AIME24	28.7	30.0	26.7	33.3	25.3	33.3	33.3	36.0
	AMC23	70.5	73.0	69.5	75.5	71.5	76.0	69.5	74.5
	GSM8K	66.7	66.8	67.2	69.2	67.0	69.1	67.3	70.8
	Math500	62.6	62.2	59.6	63.2	61.8	62.6	61.4	63.8
Scientific QA	MMLU-Pro	51.5	50.9	52.4	54.7	52.5	54.6	50.4	54.5
	GPQA	39.9	42.4	39.4	44.9	42.4	42.9	37.4	44.4
Reasoning	MuSR	52.6	53.8	52.1	54.1	52.1	52.4	52.0	56.0
	DROP	50.8	51.8	55.5	50.2	51.0	45.1	50.4	57.3
	QASC	82.1	82.9	83.6	84.1	82.2	83.3	81.4	84.4
QA	2wiki	26.2	26.4	27.0	31.6	27.2	31.4	27.1	32.5
	HotpotQA	18.1	17.3	19.1	19.7	16.9	19.6	18.1	19.9
	2wiki_RAG	36.7	33.1	33.9	37.9	32.6	33.1	33.5	41.7
QA-RAG	HotpotQA_RAG	27.1	24.5	26.0	27.3	24.7	25.2	25.7	29.2

1483 Table 13: Results of controlled comparisons for RQ2 and RQ3. We use QWEN3-8B as the base
 1484 model. This experiment is conducted on the RewardBench2 dataset. All models are trained for the
 1485 same number of steps to ensure a fair comparison. For each row within a comparison, the highest
 1486 score is in **bold**.

Task Domain	Dataset	For RQ2, RQ3.1				For RQ3.2			
		Native	RLVR @ANY	RLVR @T+F	DRM @ANY	RLVR @T+T	DRM @T+T	RLVR @F+F	DRM @F+F
Code	CodeMMLU	77.9	78.7	78.4	80.3	77.5	79.9	78.9	79.3
	CodeScope	86.5	86.8	86.2	87.4	86.9	87.6	86.7	88.3
	Cruxeval	91.6	92.2	91.9	93.0	91.5	92.6	92.1	92.5
	Execution-v2	98.5	98.7	98.7	99.0	98.3	98.5	99.0	99.0
Preference	RM-Bench	85.4	84.1	84.2	85.6	85.0	85.9	85.2	85.6
	UltraFeedback	71.3	71.8	72.9	73.2	72.4	73.2	71.7	72.2
Math	AIME24	38.0	43.3	36.7	44.7	40.7	42.7	38.7	42.0
	AMC23	72.0	74.0	69.0	79.0	73.0	80.0	74.0	76.5
	GSM8K	95.6	95.4	95.4	96.1	95.5	95.6	95.5	95.5
	Math500	73.2	74.4	72.0	75.6	73.6	75.0	72.8	75.0
Scientific QA	MMLU-Pro	65.3	64.4	61.5	71.4	65.2	71.2	64.2	68.9
	GPQA	48.0	45.5	46.0	58.1	46.0	54.5	47.0	54.5
Reasoning	MuSR	63.5	61.8	62.7	65.5	63.2	65.3	63.1	64.0
	DROP	74.7	74.2	74.2	74.9	74.9	75.3	75.2	75.4
	QASC	94.1	93.8	93.7	94.2	93.7	94.0	93.7	94.0
QA	2wiki	39.8	40.6	41.0	42.2	40.0	41.3	40.2	41.1
	HotpotQA	29.2	28.1	27.9	29.4	28.4	28.7	28.7	29.7
	2wiki_RAG	55.7	55.4	55.4	56.1	55.7	56.2	55.4	56.0
QA-RAG	HotpotQA_RAG	40.5	38.9	39.2	40.7	40.1	40.5	39.9	41.0

G.4 ABLATION STUDY

We conduct thorough ablation experiments on each supervision dimension, for each model and each training dataset, as shown in Table 18, Table 19, Table 20, Table 21, Table 22 and Table 23. Across

1512 Table 14: Results of controlled comparisons for RQ2 and RQ3. This experiment is conducted
 1513 on the HotpotQA with RAG dataset. We use LLAMA3.1-8B-INSTRUCT as the base model. All
 1514 models are trained for the same number of steps to ensure a fair comparison. For each row within a
 1515 comparison, the highest score is in **bold**.

Task Domain	Dataset	For RQ2, RQ3.1				For RQ3.2			
		Native	RLVR @ANY	RLVR @T+F	DRM @ANY	RLVR @T+T	DRM @T+T	RLVR @F+F	DRM @F+F
Code	CodeMMLU	58.8	57.2	59.5	60.5	57.6	59.4	57.2	57.4
	CodeScope	34.8	36.0	37.6	41.7	37.5	41.5	34.0	39.4
	Cruxeval	50.4	53.1	53.5	56.2	52.9	55.5	51.5	56.5
	Execution-v2	38.2	40.3	41.1	43.4	38.4	46.8	40.1	43.8
Preference	RM-Bench	56.4	59.7	56.5	62.9	59.9	60.1	58.5	61.8
	UltraFeedback	66.6	66.6	64.8	68.2	64.4	67.2	65.6	67.8
Math	AIME24	4.7	2.7	4.7	3.3	4.0	5.3	2.0	4.7
	AMC23	22.5	21.5	21.0	28.5	25.0	23.5	20.0	27.0
	GSM8K	88.8	90.0	88.8	91.5	89.4	90.2	86.7	92.1
	Math500	39.6	41.0	40.6	45.0	41.2	44.2	39.8	44.2
Scientific QA	MMLU-Pro	41.9	46.5	47.1	49.6	45.0	48.6	44.6	48.1
	GPQA	31.3	33.3	29.3	34.3	24.2	31.3	25.8	31.3
Reasoning	MuSR	48.3	48.7	49.2	53.0	49.7	50.4	49.5	49.7
	DROP	56.9	56.0	62.9	67.3	59.2	61.0	57.0	58.2
	QASC	84.4	86.9	86.0	87.5	85.3	85.2	84.4	86.3
	QA	2wiki	33.8	32.9	38.3	40.9	36.1	35.2	33.3
QA-RAG	HotpotQA	29.3	29.4	31.5	32.8	30.8	30.2	27.7	29.6
	2wiki_RAG	31.2	35.7	47.0	48.4	37.5	41.0	31.6	38.6
	HotpotQA_RAG	28.3	28.3	35.1	40.8	30.8	33.9	28.8	32.7

1537 Table 15: Results of controlled comparisons for RQ2 and RQ3. This experiment is conducted on the
 1538 HotpotQA with RAG dataset. We use R1-DISTIL-LLAMA8B as the base model. All models are
 1539 trained for the same number of steps to ensure a fair comparison. For each row within a comparison,
 1540 the highest score is in **bold**.

Task Domain	Dataset	For RQ2, RQ3.1				For RQ3.2			
		Native	RLVR @ANY	RLVR @T+F	DRM @ANY	RLVR @T+T	DRM @T+T	RLVR @F+F	DRM @F+F
Code	CodeMMLU	59.7	62.0	64.4	66.6	61.6	65.0	60.2	65.5
	CodeScope	67.4	67.0	68.3	69.7	65.0	67.6	65.6	65.6
	Cruxeval	71.9	74.6	74.6	75.4	74.0	75.8	73.4	73.8
	Execution-v2	80.8	81.2	83.1	85.6	82.9	85.2	80.4	85.0
Preference	RM-Bench	71.9	69.6	70.8	72.9	66.2	70.7	69.8	70.7
	UltraFeedback	65.2	64.6	65.4	67.0	63.3	64.8	64.3	66.6
Math	AIME24	28.7	29.3	30.0	30.7	28.0	36.7	32.0	36.0
	AMC23	70.5	67.5	70.0	80.5	70.5	81.5	72.0	78.5
	GSM8K	66.7	67.0	69.1	86.4	66.1	87.4	66.2	78.6
	Math500	62.6	58.4	59.6	67.2	61.2	67.2	58.0	66.2
Scientific QA	MMLU-Pro	51.5	51.5	52.6	54.9	53.2	53.7	50.4	55.4
	GPQA	39.9	41.4	44.9	41.4	42.4	42.9	43.4	44.4
Reasoning	MuSR	52.6	55.2	55.4	58.6	52.2	55.7	52.9	57.4
	DROP	50.8	48.6	64.3	65.4	50.7	54.3	47.1	48.6
	QASC	82.1	82.0	84.6	85.2	82.5	84.6	81.5	85.1
	QA	26.2	24.7	34.2	37.9	26.8	30.9	16.5	7.2
QA-RAG	HotpotQA	18.1	16.2	21.9	24.0	16.9	20.8	16.1	16.3
	2wiki_RAG	36.7	28.7	52.7	51.6	32.8	37.3	25.0	33.0
	HotpotQA_RAG	27.1	25.2	37.5	37.2	23.8	27.8	22.0	27.9

1562 all settings, the results show a consistent trend: no single dimension is sufficient to yield robust
 1563 improvements across diverse tasks. Combining multiple complementary dimensions produces co-
 1564 operative effects that enhance generalization and no single dimension is dominant.

1566
1567
1568
1569
1570
1571Table 16: Results of controlled comparisons for RQ2 and RQ3. This experiment is conducted on the HotpotQA with RAG dataset. We use QWEN3-8B as the base model. All models are trained for the same number of steps to ensure a fair comparison. For each row within a comparison, the highest score is in **bold**.1572
1573
1574
1575
1576
1577
1578
1579
1580
1581
1582
1583
1584
1585
1586
1587
1588
1589
1590
1591

Task Domain	Dataset	For RQ2, RQ3.1				For RQ3.2			
		Native	RLVR @ ANY	RLVR @ T+F	DRM @ ANY	RLVR @ T+T	DRM @ T+T	RLVR @ F+F	DRM @ F+F
Code	CodeMMLU	77.9	78.0	78.0	79.0	77.7	79.7	78.5	78.3
	CodeScope	86.5	87.0	87.1	87.3	86.5	86.7	87.1	87.7
	Cruxeval	91.6	91.1	92.8	92.2	92.2	92.4	92.5	91.6
	Execution-v2	98.5	98.7	98.7	98.7	98.7	98.5	98.1	98.3
Preference	RM-Bench	85.4	85.4	84.5	85.2	84.8	85.0	84.2	84.7
	UltraFeedback	71.3	72.6	73.0	72.7	72.8	72.6	71.8	73.7
Math	AIME24	38.0	42.7	40.0	47.3	40.0	46.0	40.0	44.7
	AMC23	72.0	76.0	75.5	82.5	74.0	81.0	73.0	77.0
	GSM8K	95.6	95.7	95.7	96.0	95.5	96.0	95.7	95.8
	Math500	73.2	74.0	72.8	76.4	74.6	76.8	73.4	75.6
Scientific QA	MMLU-Pro	65.3	65.4	64.1	70.4	64.8	71.4	64.0	71.2
	GPQA	48.0	46.0	46.5	56.1	49.5	59.1	46.0	55.6
Reasoning	MuSR	63.5	64.6	63.4	63.5	63.5	63.8	62.8	63.1
	DROP	74.7	73.7	75.4	74.2	75.6	73.9	74.7	74.7
	QASC	94.1	93.4	93.6	94.0	93.7	94.5	93.4	93.9
QA	2wiki	39.8	39.5	39.7	40.1	40.5	40.7	40.9	39.8
	HotpotQA	29.2	27.8	28.6	28.7	29.2	28.9	28.5	28.5
	2wiki_RAG	55.7	55.9	56.4	56.9	55.0	55.7	55.7	55.6
QA-RAG	HotpotQA_RAG	40.5	39.2	40.3	38.8	39.6	39.5	39.4	38.5

1592
1593
1594
1595
1596
1597
1598
1599Table 17: Results of on-policy GRPO training on RewardBench2. **RLVR** denotes training with answer supervision signals only. **DRM** denotes training with reasoning supervision signals only. **Combination** denotes training with their combination. For each row within a model group, the highest score is in **bold**.1600
1601
1602
1603
1604
1605
1606
1607
1608
1609
1610
1611
1612
1613
1614
1615
1616
1617
1618
1619

Task Domain	Dataset	LLaMA3.1-8B-Instruct			R1-distil-LLaMA8B			Qwen3-8B		
		RLVR	DRM	Combination	RLVR	DRM	Combination	RLVR	DRM	Combination
Code	CodeMMLU	57.0	58.0	59.0	62.4	65.1	64.0	78.0	79.1	79.2
	CodeScope	37.2	39.4	40.5	69.2	68.2	70.8	87.3	87.7	87.5
	Cruxeval	55.6	54.8	56.4	74.4	76.0	76.1	92.9	92.8	91.9
	Execution-v2	44.7	42.4	46.4	82.3	83.5	85.6	98.5	99.0	99.2
Preference	RM-Bench	59.5	57.7	60.5	73.6	65.3	69.0	85.6	72.8	83.5
	UltraFeedback	63.1	65.2	65.5	63.4	64.0	63.9	73.0	65.1	72.5
Math	AIME24	4.7	4.7	4.7	29.3	34.7	33.3	38.0	46.7	45.3
	AMC23	20.5	23.0	24.5	70.5	77.5	80.5	75.0	81.5	79.5
	GSM8K	90.7	89.6	92.3	72.5	83.1	83.0	95.1	96.1	96.0
	Math500	40.8	38.0	45.4	62.8	67.0	67.2	73.8	75.8	75.8
Scientific QA	MMLU-Pro	42.3	43.2	47.8	53.6	53.4	54.1	63.7	68.7	69.1
	GPQA	30.8	28.8	32.3	39.4	43.9	42.4	43.9	57.6	56.6
Reasoning	MuSR	47.6	52.9	52.1	53.0	53.0	52.9	63.0	63.2	64.3
	DROP	62.3	61.8	63.3	54.3	42.5	50.0	74.6	74.8	74.4
	QASC	83.3	83.5	85.1	83.8	84.6	83.5	93.4	94.1	94.2
QA	2wiki	29.5	26.3	30.6	26.7	24.4	27.6	40.6	42.2	41.4
	HotpotQA	28.6	28.1	29.1	18.5	17.3	19.5	27.7	29.5	28.6
	2wiki_RAG	34.1	33.2	34.3	36.5	24.7	29.2	56.0	55.6	55.1
QA-RAG	HotpotQA_RAG	31.0	31.4	31.9	27.1	21.7	23.9	39.3	39.6	38.9

1620
 1621
 1622
 1623 Table 18: Ablation results of single dimension supervised training LLAMA3.1-8B-INSTRUCT on
 1624 RewardBench2. Training pairs are selected from ANY subset. **DRM** means training with DRM
 1625 supervision. All training pairs are selected from ANY subset.

Task	Domain	Dataset	Native	Confidence	Coherence	Relevance	DRM
Code		CodeMMLU	58.8	57.5	58.4	55.1	59.9
		CodeScope	34.8	39.7	39.5	31.8	41.1
		Cruxeval	50.4	53.9	53.5	32.4	57.5
		Execution-v2	38.2	40.3	44.3	38.4	45.3
Preference		RM-Bench	56.4	59.2	60.8	59.1	61.0
		UltraFeedback	66.6	65.3	67.8	64.7	69.9
Math		AIME24	4.7	4.0	4.7	2.7	6.0
		AMC23	22.5	23.0	24.5	27.5	29.5
		GSM8K	88.8	83.0	89.8	89.7	91.8
		Math500	39.6	40.4	45.8	43.4	48.4
Scientific QA		MMLU-Pro	41.9	46.7	48.7	42.0	48.7
		GPQA	31.3	28.3	29.8	25.3	35.9
Reasoning		MuSR	48.3	48.7	50.8	46.2	51.7
		DROP	56.9	50.4	64.5	27.6	63.6
		QASC	84.4	84.0	86.3	77.2	87.2
QA		2wiki	33.8	34.9	32.2	29.0	35.6
		HotpotQA	29.3	29.6	30.0	26.0	31.8
QA-RAG		2wiki_RAG	31.2	28.5	36.1	31.2	39.9
		HotpotQA_RAG	28.3	27.1	33.1	27.4	34.5

1620
 1621
 1622
 1623 Table 19: Ablation results of single dimension supervised training R1-DISTIL-LLAMA8B on Re-
 1624 wardBench2. Training pairs are selected from ANY subset. **DRM** means training with DRM super-
 1625 vision. All training pairs are selected from ANY subset.

Task	Domain	Dataset	Native	Confidence	Coherence	Relevance	DRM
Code		CodeMMLU	59.7	60.2	63.9	62.8	66.3
		CodeScope	67.4	65.9	67.4	66.3	70.2
		Cruxeval	71.9	73.5	76.1	73.4	77.2
		Execution-v2	80.8	81.0	84.6	81.8	86.0
Preference		RM-Bench	71.9	71.3	70.5	68.7	74.6
		UltraFeedback	65.2	64.6	64.8	65.0	66.8
Math		AIME24	28.7	26.0	29.3	31.3	33.3
		AMC23	70.5	70.5	71.0	69.5	75.5
		GSM8K	66.7	69.7	67.8	73.2	69.2
		Math500	62.6	60.4	62.6	64.2	63.2
Scientific QA		MMLU-Pro	51.5	52.3	52.9	51.6	54.7
		GPQA	39.9	39.9	39.9	40.9	44.9
Reasoning		MuSR	52.6	51.9	54.2	52.5	54.1
		DROP	50.8	56.4	55.3	29.5	50.2
		QASC	82.1	81.6	83.2	80.8	84.1
QA		2wiki	26.2	26.6	30.5	15.0	31.6
		HotpotQA	18.1	17.8	19.1	13.6	19.7
		2wiki_RAG	36.7	39.3	39.1	20.2	37.9
QA-RAG		HotpotQA_RAG	27.1	27.1	27.9	17.9	27.3

1674
1675
1676
1677
1678
1679Table 20: Ablation results of single dimension supervised training QWEN3-8B on RewardBench2. Training pairs are selected from ANY subset. **DRM** means training with DRM supervision. All training pairs are selected from ANY subset.1680
1681
1682
1683
1684
1685
1686
1687
1688
1689
1690
1691
1692
1693
1694
1695
1696
1697
1698
1699
1700
1701
1702
1703
1704
1705
1706
1707
1708
1709
1710
1711
1712
1713
1714
1715
1716
1717
1718
1719
1720
1721
1722
1723
1724
1725
1726
1727

Task Domain	Dataset	Native	Confidence	Coherence	Relevance	DRM
Code	CodeMMLU	77.9	78.0	79.5	78.1	80.3
	CodeScope	86.5	86.8	86.3	86.7	87.4
	Cruxeval	91.6	92.9	91.8	91.9	93.0
	Execution-v2	98.5	98.1	98.3	98.3	99.0
Preference	RM-Bench	85.4	84.8	84.6	84.8	85.6
	UltraFeedback	71.3	71.1	72.0	72.0	73.2
Math	AIME24	38.0	38.7	39.3	45.3	44.7
	AMC23	72.0	72.5	77.5	79.5	79.0
	GSM8K	95.6	95.2	95.7	95.4	96.1
	Math500	73.2	72.4	74.6	74.2	75.6
Scientific QA	MMLU-Pro	65.3	62.5	70.5	70.7	71.4
	GPQA	48.0	43.9	54.0	56.1	58.1
Reasoning	MuSR	63.5	64.8	63.5	62.8	65.5
	DROP	74.7	74.4	74.0	74.6	74.9
	QASC	94.1	93.8	93.7	94.0	94.2
QA	2wiki	39.8	40.9	39.5	42.0	42.2
	HotpotQA	29.2	28.3	28.7	27.3	29.4
QA-RAG	2wiki_RAG	55.7	55.7	55.1	55.9	56.1
	HotpotQA_RAG	40.5	40.2	40.0	40.2	40.7

Table 21: Ablation results of single dimension supervised training LLAMA3.1-8B-INSTRUCT on HotpotQA with RAG. Training pairs are selected from ANY subset. **DRM** means training with DRM supervision. All training pairs are selected from ANY subset.

Task Domain	Dataset	Native	Confidence	Coherence	Relevance	DRM
Code	CodeMMLU	58.8	58.7	59.6	58.1	60.5
	CodeScope	34.8	37.6	41.8	39.3	41.7
	Cruxeval	50.4	54.5	54.0	52.5	56.2
	Execution-v2	38.2	42.8	43.6	39.7	43.4
Preference	RM-Bench	56.4	59.9	60.3	59.4	62.9
	UltraFeedback	66.6	65.6	66.4	65.3	68.2
Math	AIME24	4.7	3.3	6.7	5.3	3.3
	AMC23	22.5	22.5	26.0	19.5	28.5
	GSM8K	88.8	87.8	89.6	90.6	91.5
	Math500	39.6	41.2	46.2	41.8	45.0
Scientific QA	MMLU-Pro	41.9	46.2	47.9	46.5	49.6
	GPQA	31.3	28.3	29.3	31.8	34.3
Reasoning	MuSR	48.3	48.0	51.6	50.7	53.0
	DROP	56.9	62.3	65.3	58.9	67.3
	QASC	84.4	83.4	85.4	87.5	87.5
QA	2wiki	33.8	35.2	38.4	37.6	40.9
	HotpotQA	29.3	31.7	33.2	30.7	32.8
QA-RAG	2wiki_RAG	31.2	32.8	45.5	43.0	48.4
	HotpotQA_RAG	28.3	30.0	38.7	34.0	40.8

1728
 1729
 1730
 1731 Table 22: Ablation results of single dimension supervised training R1-DISTIL-LLAMA8B on Hot-
 1732 potQA with RAG. Training pairs are selected from ANY subset. **DRM** means training with DRM
 1733 supervision. All training pairs are selected from ANY subset.

Task	Domain	Dataset	Native	Confidence	Coherence	Relevance	DRM
Code		CodeMMLU	59.7	64.2	64.6	65.4	66.6
		CodeScope	67.4	67.3	68.5	69.2	69.7
		Cruxeval	71.9	71.9	75.2	73.1	75.4
		Execution-v2	80.8	80.8	83.5	81.8	85.6
Preference		RM-Bench	71.9	71.9	70.0	70.2	72.9
		UltraFeedback	65.2	65.0	64.0	64.4	67.0
Math		AIME24	28.7	28.0	32.7	30.7	30.7
		AMC23	70.5	64.0	79.5	77.5	80.5
		GSM8K	66.7	66.2	84.7	87.8	86.4
		Math500	62.6	58.6	65.8	65.4	67.2
Scientific QA		MMLU-Pro	51.5	52.4	51.6	53.5	54.9
		GPQA	39.9	41.4	39.9	38.4	41.4
Reasoning		MuSR	52.6	54.8	56.0	57.5	58.6
		DROP	50.8	63.0	63.0	41.0	65.4
		QASC	82.1	84.6	84.1	83.1	85.2
QA		2wiki	26.2	24.3	37.3	7.1	37.9
		HotpotQA	18.1	20.2	22.6	13.9	24.0
QA-RAG		2wiki_RAG	36.7	46.0	49.6	25.6	51.6
		HotpotQA_RAG	27.1	33.5	35.8	22.6	37.2

1752
 1753
 1754
 1755
 1756
 1757
 1758 Table 23: Ablation results of single dimension supervised training QWEN3-8B on HotpotQA with
 1759 RAG. Training pairs are selected from ANY subset. **DRM** means training with DRM supervision.
 1760 All training pairs are selected from ANY subset.

Task	Domain	Dataset	Native	Confidence	Coherence	Relevance	DRM
Code		CodeMMLU	77.9	77.3	79.0	79.2	79.0
		CodeScope	86.5	86.6	86.6	86.8	87.3
		Cruxeval	91.6	92.1	92.0	91.9	92.2
		Execution-v2	98.5	98.1	97.9	98.3	98.7
Preference		RM-Bench	85.4	85.9	84.6	83.8	85.2
		UltraFeedback	71.3	72.0	71.4	72.2	72.7
Math		AIME24	38.0	38.0	44.0	45.3	47.3
		AMC23	72.0	73.0	78.5	83.5	82.5
		GSM8K	95.6	95.5	95.4	95.7	96.0
		Math500	73.2	72.0	75.6	76.2	76.4
Scientific QA		MMLU-Pro	65.3	62.4	71.1	70.1	70.4
		GPQA	48.0	42.9	52.0	55.6	56.1
Reasoning		MuSR	63.5	63.9	62.7	63.2	63.5
		DROP	74.7	73.2	73.5	74.4	74.2
		QASC	94.1	92.7	93.7	94.0	94.0
QA		2wiki	39.8	40.9	38.9	37.5	40.1
		HotpotQA	29.2	27.2	28.5	28.2	28.7
QA-RAG		2wiki_RAG	55.7	55.5	55.5	53.9	56.9
		HotpotQA_RAG	40.5	38.7	38.1	35.4	38.8