

EFFICIENT REINFORCEMENT FINETUNING VIA ADAPTIVE CURRICULUM LEARNING

000
001
002
003
004
005 **Anonymous authors**
006 Paper under double-blind review
007
008
009
010

ABSTRACT

011 Reinforcement finetuning (RFT) has shown great potential for enhancing the math-
012 ematical reasoning capabilities of large language models (LLMs), but it is often
013 sample- and compute-inefficient, requiring extensive training. In this work, we
014 introduce ADARFT (*Adaptive Curriculum Reinforcement Finetuning*), a method
015 that significantly improves both the efficiency and final accuracy of RFT through
016 adaptive curriculum learning. ADARFT dynamically adjusts the difficulty of
017 training problems based on the model’s recent reward signals, ensuring that the
018 model consistently trains on tasks that are challenging but solvable. This adaptive
019 sampling strategy accelerates learning by maintaining an optimal difficulty range,
020 avoiding wasted computation on problems that are too easy or too hard. ADARFT
021 requires only a lightweight extension to standard RFT algorithms like Proximal
022 Policy Optimization (PPO), without modifying the reward function or model archi-
023 tecture. Experiments on competition-level math datasets—including AMC, AIME,
024 and IMO-style problems—demonstrate that ADARFT significantly improves both
025 training efficiency and reasoning performance. We evaluate ADARFT across mul-
026 tiple data distributions and model sizes, showing that it reduces training time by up
027 to $2\times$ and improves accuracy by a considerable margin, offering a more scalable
028 and effective RFT framework.
029

1 INTRODUCTION

030 Reinforcement Finetuning (RFT) has emerged as a powerful technique for aligning large language
031 models (LLMs) with task-specific goals, particularly in domains such as mathematics and code
032 generation where correctness is well defined (DeepSeek-AI et al., 2025; OpenAI et al., 2024b). By
033 optimizing a policy model with reward signals that reflect task success, RFT enables more targeted
034 learning than supervised finetuning (SFT) alone. However, despite its promise, RFT remains sample-
035 inefficient and computationally expensive. Its training involves repeated rollout generation, reward
036 computation, and policy updates—making it costly and difficult to scale (Ahmadian et al., 2024;
037 Kazemnejad et al., 2024; Li et al., 2024; Hu, 2025; Cui et al., 2025). Recent efforts to address RFT
038 inefficiency have focused on algorithmic simplification (e.g., RAFT (Dong et al., 2023), GRPO
039 (DeepSeek-AI et al., 2025), ReMax (Li et al., 2024)), and data-centric strategies (e.g., LIMO (Ye
040 et al., 2025), LIMR (Li et al., 2025)). While these approaches improve sample or compute efficiency,
041 they often introduce trade-offs: algorithmic simplifications may increase variance or limit stability,
042 and static data filtering or scoring can be brittle, computationally heavy, or model-specific. Moreover,
043 most methods’ success relies on fixed datasets or training schedules, which can be suboptimal in
044 non-uniform or imbalanced data regimes. More recently, early efforts have introduced curriculum-like
045 ideas into RFT. Staged curricula divide training into a few manually-defined phases of increasing
046 difficulty (Wen et al., 2025; Luo et al., 2025; Song et al., 2025), but these are coarse-grained and
047 lack adaptivity. Other methods use online data filtering, repeatedly rolling out and pruning training
048 samples until the model’s average reward meets a target threshold (Bae et al., 2025; Yu et al., 2025).
049 While this approach helps prevent the model from stagnating on problems that are either too easy or
050 too difficult, it is not truly adaptive and incurs significant rollout overhead.

051 To address these limitations, we propose ADARFT, a reinforcement finetuning method based on
052 adaptive curriculum learning (Bengio et al., 2009), which dynamically adjusts training set difficulty to
053 match the model’s evolving skill level. The intuition is simple: learning is most effective when tasks
are neither too easy nor too hard. ADARFT formalizes this by maintaining a target difficulty level,

which increases or decreases based on recent reward feedback. At each step, the model is trained on examples closest to this target, promoting a steady progression through solvable yet challenging tasks. The full algorithm is outlined in Algorithm 1. Unlike prior work that relies on fixed stages, repeated rollouts, or model-specific data processing, ADARFT is lightweight, general, and model-agnostic. It can be directly applied on top of any standard reinforcement learning (RL) algorithms like Proximal Policy Optimization (PPO) (Schulman et al., 2017b). We evaluate ADARFT on a dataset spanning a wide range of competition-level math problems, including AMC, AIME, and IMO-style questions. Across multiple training distributions and two model sizes, ADARFT significantly improves both training efficiency and final performance. Gains are especially notable in imbalanced data regimes, where static sampling often fails. ADARFT can reduce training time by up to 2 \times , offering a practical and scalable path to more efficient RFT in structured reasoning tasks.

2 RELATED WORK

Efficient Reinforcement Finetuning. Most RFT pipelines build on Proximal Policy Optimization (PPO) (Schulman et al., 2017b), with recent variants like RAFT (Dong et al., 2023), ReMax (Li et al., 2024), GRPO (DeepSeek-AI et al., 2025), and REINFORCE++ (Hu, 2025), aiming to reduce computational overhead by simplifying RL components. While effective, these methods often trade off stability or sample efficiency. In parallel, data-centric strategies have emerged as promising alternatives for efficient finetuning. LIMO (Ye et al., 2025) and s1 (Muennighoff et al., 2025) show that small, carefully selected supervised datasets can yield strong downstream performance, but their success hinges on manual curation, prompt engineering, and careful dataset construction, which may not generalize across tasks or models. LIMR (Li et al., 2025) and Wang et al. (2025) proposes scoring training examples based on their estimated learning impact, enabling selective finetuning with fewer samples. Yet, computing these scores requires a full training run, and the scores must be recomputed for each new model, limiting practicality and scalability. Moreover, reducing the number of training samples does not inherently translate to improved efficiency. Models still require a comparable number of optimization steps and wall-clock time to converge. In contrast, ADARFT introduces a lightweight, model-agnostic curriculum learning strategy that dynamically adjusts task difficulty based on reward feedback. This allows continuous adaptation to the model’s capabilities, improving convergence speed and final accuracy without modifying the RL algorithm or requiring manual data curation.

Curriculum Learning for RL. Curriculum learning (CL) structures training by presenting tasks in an organized progression, typically from easy to hard, to enhance learning efficiency and generalization (Bengio et al., 2009). In RL, CL methods include task sorting by difficulty (Zaremba & Sutskever, 2015; Justesen et al., 2018; Wang et al., 2019), teacher-student frameworks that adaptively select tasks based on learning progress (Matiisen et al., 2017; Portelas et al., 2019), and self-play approaches that induce automatic curricula through agent competition (Sukhbaatar et al., 2018; Zhao et al., 2025). Other strategies use intermediate-goal generation in sparse-reward settings (Florensa et al., 2018), unsupervised skill discovery (Jabri et al., 2019), or knowledge transfer via progressive networks and imitation (Czarnecki et al., 2018; Rusu et al., 2022). While CL is well-studied in classical RL, its application to RFT of LLMs is still limited. Existing methods typically use staged training with hand-designed difficulty tiers (Wen et al., 2025; Luo et al., 2025; Song et al., 2025), or online filtering schemes that repeatedly sample and discard data until rewards reach a target range (Bae et al., 2025; Yu et al., 2025). These methods either lack adaptability or introduce significant computational overhead due to repeated rollouts. In contrast, ADARFT is among the first truly adaptive curriculum learning approaches for RFT: it continuously adjusts task difficulty based on the model’s reward signal, enabling efficient, scalable training without fixed schedules or repeated rollouts.

3 ADARFT

We aim to improve the performance of a policy model π_θ for solving mathematical problems through adaptive curriculum learning. Fine-tuning on problems that are too easy or too hard leads to poor learning outcomes. Instead, the model should be trained on problems whose difficulty is close to the model’s current capability. We frame this as an adaptive curriculum learning problem and propose

108 ADARFT, which adaptively adjusts the target difficulty to keep training problems within a suitable
 109 difficulty range. ADARFT is compatible with a variety of RL algorithms (e.g, GRPO, PPO); in this
 110 work, we instantiate it with PPO and refer to this variant as ADARFT (PPO).

111 Let D be a dataset of mathematical problems, each annotated with a precomputed difficulty score d_i .
 112 The score can be either human-annotated or model-estimated. The objective is to train a policy π_θ
 113 that improves its problem-solving ability by dynamically adjusting the training curriculum according
 114 to the model’s current performance. Our proposed algorithm, ADARFT, is shown in Algorithm 1.
 115

Algorithm 1 ADARFT – Adaptive Curriculum Reinforcement Finetuning

```

1: Input: Data source  $D$  with difficulty scores  $\{d_i\}$ , policy model  $\pi_\theta$ , reward function  $R(\cdot, \cdot)$ , batch size  $B$ ,  

  initial target difficulty  $T$ , step size  $\eta$ , sensitivity  $\alpha$ , target reward  $\beta$ , difficulty bounds  $d_{\min}, d_{\max}$ 
2: Select RL algorithm  $\mathcal{A}$  (e.g., PPO, GRPO, REINFORCE++)
3: while training is not finished do
4:   Compute absolute differences from target difficulty:  $\Delta_i = |d_i - T| \quad \forall i \in \{1, \dots, |D|\}$ 
5:   Sort and select top  $B$  samples closest to target difficulty:  $X \leftarrow \{s_1, s_2, \dots, s_B\}$ 
6:   Generate responses using policy model:  $G = \pi_\theta(X)$ 
7:   Compute average reward:  $R_{avg} \leftarrow \frac{1}{|X|} \sum_{i=1}^{|X|} R(X_i, G_i)$ 
8:   Update policy:  $\pi_\theta \leftarrow \mathcal{A}(\pi_\theta, X, G, R)$ 
9:   Update and clip target difficulty:  $T' \leftarrow \text{clip}(T + \eta \cdot \tanh(\alpha \cdot (R_{avg} - \beta)), d_{\min}, d_{\max})$ 
10:  Update sampler:  $T \leftarrow T'$ 
11: end while

```

3.1 DYNAMIC CURRICULUM SAMPLING

To construct an adaptive curriculum, we define a target difficulty T , which represents the current target difficulty level for training (more in § 3.3). ADARFT dynamically adjusts T based on the model’s reward signal to maintain an optimal difficulty level for learning. At each step, the algorithm computes the absolute difference between the target difficulty and the difficulty of each problem in the dataset (Alg. 1, line 4): $\Delta_i = |d_i - T|$ for all $i \in [1, |D|]$. The batch of training problems is formed by selecting the B problems with the smallest values of Δ_i (Alg. 1, line 5), producing a batch: $X = \{s_1, s_2, \dots, s_B\}$. This ensures that the selected problems are closest to the model’s current target difficulty, focusing the learning process on problems that are neither too easy nor too hard.

3.2 POLICY UPDATE

The selected batch X is used to train the policy model π_θ , which generates responses: $G = \pi_\theta(X)$. A reward signal is computed based on the correctness of the model’s output (Alg. 1, line 7): $R_i = 1$ if the response is correct, and $R_i = 0$ if the response is incorrect. The average reward over the batch is computed as (Alg. 1, line 7): $R_{avg} = \frac{1}{|X|} \sum_{i=1}^{|X|} R(X_i, G_i)$. The policy can then be updated using a reinforcement learning algorithm \mathcal{A} such as PPO, GRPO, or REINFORCE++ (Alg. 1, line 8): $\pi_\theta \leftarrow \mathcal{A}(\pi_\theta, X, G, R)$.

3.3 TARGET DIFFICULTY UPDATE

To adapt the curriculum dynamically, the target difficulty is updated based on the average reward. If the model performs well on the current difficulty level (high reward), the target difficulty increases, making the training problems harder. Conversely, if the model performs poorly, the target difficulty decreases. This dynamic update mechanism lies at the core of ADARFT’s curriculum adaptation strategy. The update rule (Alg. 1, line 9) is defined as:

$$T' = \text{clip}(T + \eta \cdot \tanh(\alpha \cdot (R_{avg} - \beta)), d_{\min}, d_{\max})$$

Here, η, α, β are hyperparameters: η is the step size for adjusting the target difficulty, α controls the sensitivity of the update, and β is the target reward level, representing the desired success rate. The \tanh function ensures smooth updates and prevents large jumps in difficulty by saturating for large deviations, while the “clip” function constrains the target difficulty within the valid range

[d_{\min}, d_{\max}]. These bounds can be manually specified or automatically derived from the training set, for example by taking the minimum and maximum of the difficulty scores $\{d_i\}$. Intuition and guidance for selecting these hyperparameters are discussed in Section 3.4 and 4.3.

3.4 THEORETICAL JUSTIFICATION FOR TARGET REWARD β

A key component of ADARFT is its adaptive curriculum mechanism, which steers training toward a target reward level β . Intuitively, we aim to train on examples that are neither trivially easy nor prohibitively hard. In this light, setting $\beta = 0.5$, corresponding to a success rate of roughly 50%, naturally aligns with this goal. This section formalizes that intuition by analyzing the relationship between reward variance and learnability in RFT with binary rewards.

In entropy-regularized reinforcement learning, the optimal policy π^* can be expressed relative to a reference policy π_{init} as (Korbak et al., 2022; Go et al., 2023; Rafailov et al., 2023):

$$\pi^*(y | x) = Z(x)\pi_{\text{init}}(y | x) \exp\left(\frac{1}{\tau}r(x, y)\right) \quad (1)$$

where τ is the inverse temperature parameter controlling entropy regularization, and $Z(x)$ is the partition function that normalizes the action probability. The corresponding optimal value function and the partition function is given by (Schulman et al., 2017a; Richemond et al., 2024):

$$V^*(x) := \tau \log \mathbb{E}_{y \sim \pi_{\text{init}}(\cdot | x)} \left[\exp\left(\frac{1}{\tau}r(x, y)\right) \right] \quad \text{and} \quad Z(x) = \exp\left(\frac{1}{\tau}V^*(x)\right) \quad (2)$$

We can then take the expectation of the log-ratio between the optimal policy and the initial policy with respect to $y \sim \pi_{\text{init}}(\cdot | x)$, leading to (Haarnoja et al., 2017; Schulman et al., 2017a):

$$\mathbb{E}_{y \sim \pi_{\text{init}}(\cdot | x)} \left[\log \frac{\pi^*(y | x)}{\pi_{\text{init}}(y | x)} \right] = \frac{1}{\tau} \mathbb{E}_{\pi_{\text{init}}} [r(x, y)] - \frac{1}{\tau} V^*(x) \quad (3)$$

Since the left-hand side can be interpreted as the negative reverse KL divergence between π_{init} and π^* (Rafailov et al., 2024), Bae et al. (2025) show that when the reward $r(x, y)$ with $y \sim \pi_{\text{init}}(\cdot | x)$ is Bernoulli, the KL divergence is lower-bounded by the reward variance:

$$D_{\text{KL}}(\pi_{\text{init}} \| \pi^*) \geq \frac{p(x)(1 - p(x))}{2\tau^2} \quad (4)$$

where $p(x)$ is the model’s success rate on prompt x . This implies that the lower bound on the KL divergence, and consequently the gradient magnitude during policy updates, is proportional to the reward variance, which is maximized when $p(x) = 0.5$. In other words, training on prompts that the model succeeds on roughly half the time may yield the strongest learning signal. In Section 5 and Appendix 5.3, we conduct an ablation study by varying the target reward β , demonstrating that setting $\beta = 0.5$ consistently leads to the best performance, supporting the hypothesis that training on prompts with a success rate near 50% provides the most informative learning signal.

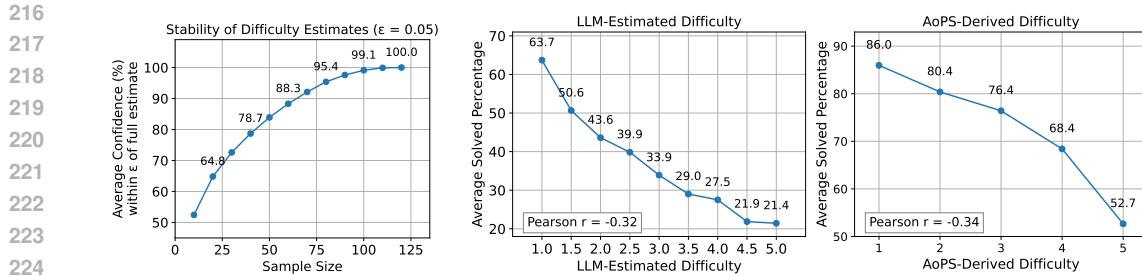
4 EXPERIMENTS

4.1 DIFFICULTY ESTIMATION

Accurate estimation of problem difficulty is critical for ADARFT. For difficulty estimation, we select the Qwen 2.5 MATH 7B model (Qwen et al., 2025) because it demonstrates a balanced solving ability. A model that is too strong (e.g., OpenAI o1 (OpenAI et al., 2024b), DeepSeek R1 (DeepSeek-AI et al., 2025)) would solve most problems on the first attempt, leading to poor discrimination between easy and hard problems. Conversely, a model that is too weak (e.g., LLaMA 3.3 1B (Grattafiori et al., 2024)) would fail to solve most problems even after multiple attempts, limiting the signal required for curriculum adaptation. For each problem, the difficulty score is computed as:

$$d_i = 100 \times \left(1 - \frac{\text{number of successful attempts on problem } i}{n} \right)$$

where n is the number of attempts per problem. In our setup, we use $n = 128$.



(a) Average confidence that subsampled difficulty estimates fall within ± 0.05 of the full-sample estimate. (b) Correlation between average solved percentage and two types of difficulty labels: (left) LLM-estimated difficulty and (right) AoPS-derived difficulty levels.

Figure 1: Evaluation of difficulty estimation: (a) Stability of difficulty scores under subsampling of model rollouts; (b) Correlation between labeled difficulty levels and average solved percentage.

To evaluate the stability of our difficulty estimation process, we simulate how confidence varies with different numbers of samples. For each problem, we treat the full set of 128 rollouts as the ground-truth difficulty estimate and compute how often sub-sampled estimates fall within a tolerance of $\epsilon = 0.05$. Specifically, we run 10 random sampling trials per sample size and average the confidence across all problems in the dataset. As shown in Figure 1a, even with as few as 64 samples, the estimated difficulty remains within ± 0.05 of the full estimate over 90% of the time. With just 40 samples, the confidence remains around 80%. These results indicate that accurate and robust difficulty estimation can be achieved with significantly fewer rollouts, reducing the computational burden of large-scale curriculum construction.

To further validate the reliability of our difficulty estimates, we examined their alignment with the difficulty levels provided in the MATH dataset. The MATH dataset comprises 12,500 competition-level mathematics problems sourced from contests such as the American Mathematics Competitions (AMC) and the American Invitational Mathematics Examination (AIME). Each problem is categorized into one of five difficulty levels, following the classification system used by the Art of Problem Solving (AoPS) community.¹ In this system, level 1 denotes the easiest problems, while level 5 represents the most difficult. As shown in Figure 1b, there is a clear downward trend in the average solve rate as the labeled difficulty level increases, ranging from 86.0% at level 1 to 52.7% at level 5. Specifically, the AoPS-derived difficulty levels yield a Pearson correlation of $r = -0.34$ ($p < 0.05$) with model success rates. This negative correlation indicates that the model’s empirical performance aligns well with the intended difficulty stratification, reinforcing the utility of both the labeled difficulty levels and our estimation approach in guiding curriculum learning. To further streamline the difficulty estimation process, we also prompted GPT-4o (gpt-4o-0806) (OpenAI et al., 2024a) to assign difficulty levels to the DeepScaleR dataset based on the AoPS rubric. Each problem was presented to GPT-4o with a request to rate its difficulty according to AoPS guidelines (the full prompt is shown in Appendix B.3). This approach provides a lightweight and scalable alternative to rollout-based estimation. As shown in Figure 1b, GPT-4o’s difficulty ratings also correlate well with the model success rates, with a Pearson correlation of $r = -0.32$ ($p < 0.05$), making it a practical proxy for curriculum scheduling when computational resources are constrained.

4.2 DATASET

We use the DeepScaleR dataset (Luo et al., 2025) as the training set. DeepScaleR compiles problems from multiple sources, including AIME from 1984 to 2023 and AMC prior to 2023. The dataset also includes problems from the Omni-MATH (Gao et al., 2024) and Still datasets (Team, 2025), which feature problems from various national and international math competitions. This results in a diverse and challenging training set, covering a wide range of mathematical domains and difficulty levels.

¹https://artofproblemsolving.com/wiki/index.php/AoPS_Wiki:Competition_ratings

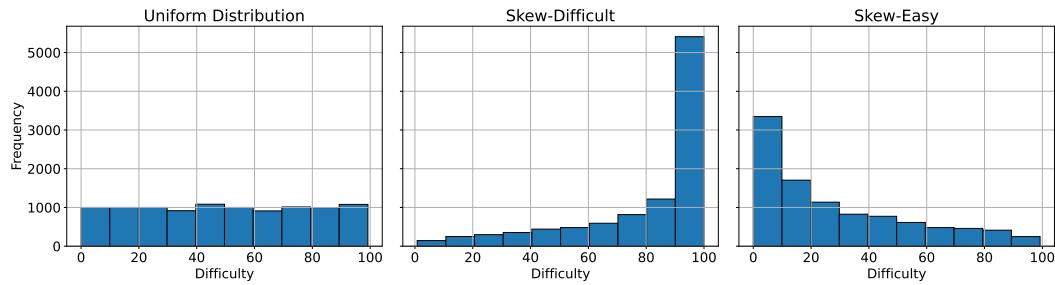


Figure 2: Difficulty distribution for different training sets: Uniform, Skew-Difficult, and Skew-Easy. Each training set contains 10,000 samples.

In practice, we do not have control over the exact difficulty distribution of the data collected for training. This motivates our investigation into how different difficulty distributions influence ADARFT. To this end, we construct three distinct distributions from the DeepScaleR dataset. The first is a *skew-difficult* distribution, where most problems are challenging. The second is a *skew-easy* distribution, where most problems are relatively easy. The third is a *uniform* distribution, where problems are evenly balanced across all difficulty levels, ensuring a consistent representation of easy, moderate, and hard problems. Each of these three distributions includes 10,000 samples. The data distribution for each setting is shown in Figure 2.

For evaluation, we use six benchmark datasets to assess the model’s performance across different levels of difficulty and mathematical reasoning. The first benchmark, MATH 500 (Lightman et al., 2023), is a subset of the MATH dataset (Hendrycks et al., 2021) containing 500 representative problems designed to test a model’s general mathematical capability. GSM8K (Cobbe et al., 2021) is a set of grade-school math problems. OlympiadBench (He et al., 2024) includes a collection of problems from Olympiad-level mathematics and physics competitions. Minerva Math (Lewkowycz et al., 2022) is a curated set of undergraduate-level math problems that assess complex mathematical reasoning and symbolic manipulation. AMC 23 and AIME 24 include problems from the 2023 American Mathematics Competitions and the 2024 American Invitational Mathematics Examination, respectively. Since AMC 23 contains only 40 problems and AIME 24 only 30, we report accuracy as the average over 8 sampled responses per problem to ensure stable estimates. Together, these datasets span elementary, high school, and advanced competition-level math, providing a comprehensive evaluation of the model’s reasoning abilities.

4.3 TRAINING SETUP

We trained two models on the three difficulty-based distributions of the DeepScaleR dataset described in Section 4.2: Qwen 2.5 7B and Qwen 2.5 MATH 1.5B. This setup allows us to evaluate the effectiveness of ADARFT on models with different initial performance levels when exposed to skew-difficult, skew-easy, and uniform problem distributions. All models were trained using four different approaches: (1) the standard PPO algorithm, (2) ADARFT (PPO), our method that integrates adaptive curriculum learning with PPO (see Section 3), (3) PPO with filtered data, a baseline that trains PPO on data filtered by pass@ k accuracy, and (4) PPO with a fixed curriculum schedule.

For the data filtering baseline (3), following prior work (Bae et al., 2025; Hu et al., 2025; Zyphra, 2025), we first run a pass@40 analysis for each combination of model and data distribution. We then discard examples that are either too easy or too hard, removing all problems with solved rates $\leq 10\%$ or $\geq 90\%$. This restricts training to problems of intermediate difficulty. However, this procedure removes a large fraction of the data, including many potentially informative examples. In addition, because difficulty is defined using pass@ k metrics, the filtering must be recomputed whenever the model or the data distribution changes.

For the fixed curriculum baseline (4), we follow the approach of prior work (Parashar et al., 2025; Team et al., 2025). In this setting, the difficulty of sampled problems follows a predetermined schedule that increases linearly over training steps. Suppose training runs for n total steps. We then

324 define a target difficulty $T(s)$ at step s by
 325

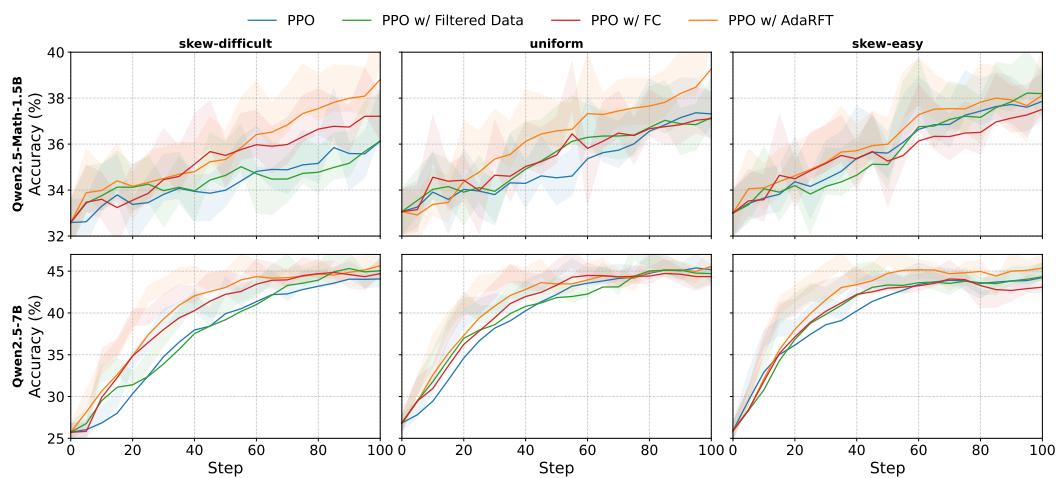
$$326 \quad 327 \quad 328 \quad T(s) = T_{\min} + \frac{T_{\max} - T_{\min}}{n} \cdot s,$$

329 and at each step we sample problems whose estimated difficulty matches this target. Unlike ADARFT,
 330 this schedule increases difficulty at a fixed rate regardless of how quickly or slowly the model learns.

331 The training batch size was set to $B = 1024$, with the target reward β set to 0.5 to promote learning
 332 at a balanced success rate. The sensitivity parameter α and step size η were tuned using a validation
 333 set to ensure stable curriculum updates. We set $\alpha = 2$, $\eta = 50$, and the initial target difficulty
 334 $T = 0$. The step size η acts as a scaling factor between the reward signal and the difficulty metric.
 335 Since the difficulty metric ranges from 0 to 100 and the reward ranges from 0 to 1, a target reward
 336 $\beta = 0.5$ implies that the maximum reasonable adjustment to the difficulty metric should be around
 337 50. Therefore, we set $\eta = 50$ to scale the reward signal appropriately to the difficulty range. The
 338 sensitivity parameter $\alpha = 2$ controls the slope of the tanh function. Setting α to 2 makes the tanh
 339 function behave approximately linearly when the difference between the average reward and the
 340 target reward is small. The intuition behind using the tanh function is that when the average reward
 341 is close to the target reward, a roughly linear adjustment is appropriate. However, when the average
 342 reward deviates significantly from the target reward, linear adjustments may be too large, leading to
 343 instability. The tanh function smooths out these adjustments, allowing for more controlled changes
 344 when the difference is large while maintaining sensitivity when the difference is small. Both models
 345 were trained on 8 A100 GPUs for approximately 100 steps. The implementation details can be found
 346 in Appendix B.

347 5 RESULTS AND ANALYSIS

350 We evaluate the performance of standard PPO and ADARFT (PPO) across multiple training setups
 351 and two model sizes: Qwen 2.5 MATH 1.5B and Qwen 2.5 7B. Figure 3 presents the learning curves
 352 averaged across six benchmarks, while Table 1 and 2 provide a detailed breakdown of accuracy and
 353 training efficiency. On average, models trained with ADARFT (PPO) outperform their PPO-only
 354 counterparts in both final accuracy and training efficiency. This improvement is particularly notable
 355 in non-uniform data distributions, where curriculum adaptation is most beneficial.



373 Figure 3: Performance comparison of PPO, PPO with filtered data, PPO with fixed curriculum (PPO
 374 w/ FC), and ADARFT (PPO) across different setups (uniform, skew-easy, skew-difficult). Accuracy
 375 is the average of MATH 500, GSM8K, AIME 24, AMC 23, OlympiadBench, and Minerva Math.
 376 Compared with baselines, ADARFT improves both the accuracy and training efficiency. For clarity,
 377 curves are exponentially smoothed.

378
379

5.1 TRAINING EFFICIENCY

380 As shown in Figure 3 and Table 1, models trained with ADARFT consistently require fewer training
 381 steps to match the performance of those trained with standard PPO, PPO on filtered data, and PPO
 382 with a fixed curriculum schedule. Specifically, we report how many additional steps are needed
 383 for PPO variants to match the performance of ADARFT at step 60 for Qwen 2.5 Math 1.5B, and
 384 step 40 for Qwen 2.5 7B. Because models are evaluated only every 5 training steps, we apply
 385 exponential smoothing with a smoothing parameter of 0.3 to the accuracy curves to reduce variance
 386 and obtain stable estimates of performance over time. The shaded areas in Figure 3 represent the raw,
 387 unsmoothed accuracy $\pm 1\%$, offering a visual cue for the typical fluctuation in evaluation accuracy.
 388 For Qwen 2.5 Math 1.5B, standard PPO requires 43 extra steps (+71.7%) in the skew-difficult setting
 389 and 34 steps (+56.7%) in the uniform setting to match ADARFT’s performance. PPO with filtered
 390 training data requires even more: +49 steps (81.7%) and +52 steps (86.7%) in the respective settings.
 391 In the skew-easy scenario, PPO requires +16 steps (26.7%), while PPO with filtered data needs +21
 392 steps (35.0%) to catch up to ADARFT. The efficiency gains remain significant with the larger Qwen
 393 2.5 7B model. In the skew-difficult setting, PPO and PPO with filtered data require +24 steps (60.0%)
 394 and +25 steps (62.5%), respectively. **PPO with a fixed curriculum schedule also follows this trend**
 395 **suggesting that while fixed curricula can modestly improve training efficiency, their inability to adapt**
 396 **the the model’s evolving learning dynamics limits their convergence speed relative to ADARFT.**

397 In addition to improved sample efficiency, ADARFT also achieves faster average training time per
 398 step across nearly all settings, as reported in Table 1. Though PPO with filtered data can sometimes
 399 offer marginal gains in per-step time (e.g., skew-easy setups), it still falls behind in convergence speed.
 400 This is largely due to the fact that easier problems require fewer tokens to solve. For example, an
 401 arithmetic reasoning question from GSM8K might require only around 200 tokens for Qwen 2.5 7B to
 402 reach a correct answer, whereas a competition-level math problem from AIME could require around
 403 2000 tokens, a $10\times$ difference in rollout length. The total token length affects multiple components
 404 of the training step, including the rollout itself and the subsequent PPO update. While PPO update
 405 time does not scale linearly with sequence length due to batching and attention computation patterns,
 406 longer sequences still incur higher compute costs. As a result, curriculum learning’s tendency to
 407 prioritize shorter, easier problems early in training leads to shorter sequences on average, reducing
 408 per-step compute and improving overall training throughput. These results underscore that ADARFT
 409 is both sample-efficient and compute-efficient, delivering faster and more cost-effective training.
 410

411 Table 1: Average time per step (in seconds) at step 100 and extra steps required to match ADARFT’s
 412 performance at step 60 (for Qwen 2.5 Math 1.5B) or step 40 (for Qwen 2.5 7B), across different
 413 setups and methods.

| 413 | Model | Setup | Method | Avg Step Time (s) | Extra Steps (%) | Extra Steps | Extra Time (s) |
|-----|-----------------|---------------------------|-----------------|-------------------|-----------------|-------------|----------------|
| 414 | Qwen2.5 1.5B | skew-difficult uniform | ADARFT | 122.24 | 0.0% | +0 | 0.00 |
| 415 | | | PPO | 132.95 | 71.7% | +43 | 5716.85 |
| 416 | | | PPO (w/ Filter) | 128.20 | 81.7% | +49 | 6281.80 |
| 417 | | | PPO (w/ FC) | 130.91 | 26.7% | +16 | 2094.56 |
| 418 | Qwen2.5 7B | skew-easy | ADARFT | 121.31 | 0.0% | +0 | 0.00 |
| 419 | | | PPO | 126.82 | 56.7% | +34 | 4311.88 |
| 420 | | | PPO (w/ Filter) | 126.35 | 86.7% | +52 | 6570.20 |
| 421 | | | PPO (w/ FC) | 126.40 | 80.0% | +48 | 6067.20 |
| 422 | Qwen2.5 7B | skew-difficult uniform | ADARFT | 120.52 | 0.0% | +0 | 0.00 |
| 423 | | | PPO | 121.15 | 26.7% | +16 | 1938.40 |
| 424 | | | PPO (w/ Filter) | 115.12 | 35.0% | +21 | 2417.52 |
| 425 | | | PPO (w/ FC) | 121.99 | 58.3% | +35 | 4269.65 |
| 426 | Qwen2.5 7B | skew-difficult uniform | ADARFT | 239.92 | 0.0% | +0 | 0.00 |
| 427 | | | PPO | 246.21 | 60.0% | +24 | 5909.04 |
| 428 | | | PPO (w/ Filter) | 254.22 | 62.5% | +25 | 6355.50 |
| 429 | | | PPO (w/ FC) | 243.12 | 22.5% | +9 | 2188.08 |
| 430 | Qwen2.5 7B | skew-easy | ADARFT | 234.16 | 0.0% | +0 | 0.00 |
| 431 | | | PPO | 243.82 | 32.5% | +13 | 3169.66 |
| 432 | | | PPO (w/ Filter) | 263.11 | 57.5% | +23 | 6051.53 |
| 433 | | | PPO (w/ FC) | 240.62 | 17.5% | +7 | 1684.34 |
| 434 | Qwen2.5 7B | skew-easy | ADARFT | 247.44 | 0.0% | +0 | 0.00 |
| 435 | | | PPO | 235.27 | 50.0% | +20 | 4705.40 |
| 436 | | | PPO (w/ Filter) | 233.13 | 42.5% | +17 | 3963.21 |
| 437 | | | PPO (w/ FC) | 240.66 | 57.5% | +23 | 5535.18 |

432 5.2 MODEL PERFORMANCE
433

434 In addition to improving efficiency, ADARFT (PPO) also improves the final model performance.
 435 As shown in Table 2, at the end of training (step 100), ADARFT yields consistent improvements in
 436 final accuracy across all configurations. The reported averages reflect accuracy across six diverse
 437 benchmarks: GSM8K, MATH 500, OlympiadBench, Minerva Math, AMC 23, and AIME 24.
 438 On skew-difficult data, the Qwen 2.5 Math 1.5B model improves from 37.41% (PPO) to 40.48%
 439 (ADARFT (PPO)), a gain of over 3 percentage points in average accuracy. Similar improvements
 440 appear in the uniform setting, where ADARFT (PPO) reaches 41.11%, compared to 37.20% with
 441 PPO. Even on skew-easy data, where the baseline performs well, curriculum learning still improves
 442 performance, reaching 39.18% versus 38.46%. For the larger Qwen 2.5 7B model, final accuracy
 443 gains are also consistent, though slightly more modest. In the skew-difficult setting, ADARFT (PPO)
 444 improves from 44.17% to 46.83%. In the uniform setting, accuracy rises from 44.70% to 46.92%,
 445 and in the skew-easy case, from 45.07% to 45.94%. These results show that ADARFT is effective
 446 even for stronger models, enhancing both stability and peak performance.
 447

448 Table 2: Accuracy (%) at step 100 for every model, setup, and benchmark. ADARFT in this table
 449 refers to ADARFT instantiated with PPO, i.e., ADARFT (PPO).

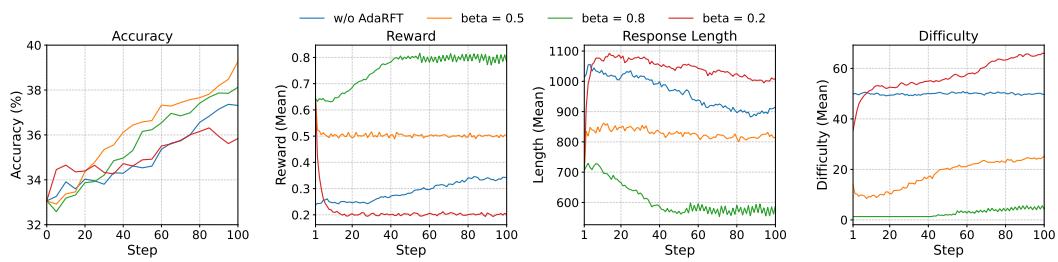
| 450 Model | 451 Setup | 452 Method | 453 GSM8K | 454 MATH 500 | 455 Olympiad Bench | 456 Minerva Math | 457 AMC 23 (Avg@8) | 458 AIME 24 (Avg@8) | 459 Average | | | |
|--|--|---|--------------------------------------|--------------------|-----------------------|---------------------|-----------------------|------------------------|------------------|------------------|------------------|------------------|
| 450 451 452 453 454 455 456 457 458 459 460 461 462 463 464 465 466 467 | 452 453 454 455 456 457 458 459 460 461 462 463 464 465 466 467 | 453 454 455 456 457 458 459 460 461 462 463 464 465 466 467 | 453 ADARFT | 454 74.00 | 455 66.40 | 456 20.36 | 457 15.07 | 458 55.00 | 459 12.08 | 460 40.48 | | |
| | | | 453 PPO | 454 69.67 | 455 64.60 | 456 20.65 | 457 12.87 | 458 47.50 | 459 9.17 | 460 37.41 | | |
| | | | 453 PPO (w/ Filter) | 454 71.65 | 455 62.40 | 456 20.06 | 457 15.07 | 458 45.00 | 459 9.17 | 460 37.22 | | |
| | | | 453 PPO (w/ FC) | 454 72.55 | 455 66.40 | 456 20.95 | 457 14.34 | 458 45.00 | 459 4.06 | 460 37.22 | | |
| | 456 457 458 459 460 461 462 463 464 465 466 467 | 456 457 458 459 460 461 462 463 464 465 466 467 | 456 uniform | 457 ADARFT | 458 74.53 | 459 66.20 | 460 21.99 | 461 14.34 | 462 57.50 | 463 12.08 | 464 41.11 | |
| | | | 456 PPO | 457 71.95 | 458 65.20 | 459 21.10 | 460 15.81 | 461 42.50 | 462 6.67 | 463 37.20 | | |
| | | | 456 PPO (w/ Filter) | 457 72.63 | 458 65.80 | 459 20.21 | 460 13.60 | 461 45.00 | 462 10.00 | 463 37.87 | | |
| | | | 456 PPO (w/ FC) | 457 71.95 | 458 65.60 | 459 19.61 | 460 14.71 | 461 42.50 | 462 9.27 | 463 37.27 | | |
| | 464 465 466 467 | 464 465 466 467 | 464 skew-easy | 465 ADARFT | 466 73.62 | 467 66.20 | 468 19.91 | 469 13.97 | 470 55.00 | 471 9.17 | | |
| | | | 464 PPO | 465 72.71 | 466 67.40 | 467 19.17 | 468 13.97 | 469 45.00 | 470 12.50 | 471 38.46 | | |
| | | | 464 PPO (w/ Filter) | 465 74.75 | 466 65.20 | 467 20.36 | 468 13.60 | 469 45.00 | 470 10.00 | 471 38.15 | | |
| | | | 464 PPO (w/ FC) | 465 72.40 | 466 66.20 | 467 20.06 | 468 13.24 | 469 50.00 | 470 6.67 | 471 38.09 | | |
| | 472 473 474 475 476 477 | 472 473 474 475 476 477 | 472 Qwen 2.5 473 Math 474 1.5B | 475 skew-difficult | 476 ADARFT | 477 90.98 | 478 71.40 | 479 25.85 | 480 22.43 | 481 52.50 | 482 15.83 | 483 46.83 |
| | | | 472 PPO | 473 89.69 | 474 71.20 | 475 23.33 | 476 23.53 | 477 50.00 | 478 11.25 | 479 44.17 | | |
| | | | 472 PPO (w/ Filter) | 473 88.48 | 474 72.20 | 475 24.37 | 476 25.00 | 477 50.00 | 478 12.08 | 479 45.35 | | |
| | | | 472 PPO (w/ FC) | 473 89.92 | 474 72.00 | 475 24.52 | 476 25.37 | 477 47.50 | 478 13.96 | 479 45.54 | | |
| | 480 481 482 483 484 485 | 480 481 482 483 484 485 | 480 Qwen 2.5 481 7B | 482 uniform | 483 ADARFT | 484 90.14 | 485 72.60 | 486 24.96 | 487 24.26 | 488 55.00 | 489 14.58 | 490 46.92 |
| | | | 482 PPO | 483 89.31 | 484 72.40 | 485 23.63 | 486 25.37 | 487 42.50 | 488 15.00 | 489 44.70 | | |
| | | | 482 PPO (w/ Filter) | 483 89.08 | 484 74.40 | 485 23.18 | 486 22.43 | 487 45.00 | 488 13.33 | 489 44.57 | | |
| | | | 482 PPO (w/ FC) | 483 89.84 | 484 72.40 | 485 23.92 | 486 23.90 | 487 42.50 | 488 13.33 | 489 44.32 | | |
| | 491 492 493 494 495 496 497 | 491 492 493 494 495 496 497 | 491 Qwen 2.5 492 7B | 493 skew-easy | 494 ADARFT | 495 90.14 | 496 72.60 | 497 25.56 | 498 23.16 | 499 50.00 | 500 14.17 | 501 45.94 |
| | | | 491 PPO | 492 89.39 | 493 73.60 | 494 23.33 | 495 24.26 | 496 47.50 | 497 13.33 | 498 45.07 | | |
| | | | 491 PPO (w/ Filter) | 492 89.31 | 493 71.60 | 494 24.22 | 495 23.90 | 496 47.50 | 497 13.33 | 498 44.98 | | |
| | | | 491 PPO (w/ FC) | 492 89.46 | 493 72.20 | 494 24.37 | 495 23.90 | 496 37.50 | 497 13.33 | 498 43.46 | | |

498 5.3 ABLATION ON TARGET REWARD β

499 To better understand the role of the target reward β in ADARFT, we perform an ablation study varying
 500 β in the target difficulty update rule. Recall that β controls the target average reward the model is
 501 expected to achieve and implicitly steers the curriculum: lower values prioritize easier problems,
 502 while higher values shift the curriculum toward more challenging samples. We train a Qwen 2.5 Math
 503 1.5B model on the uniform data distribution with ADARFT (PPO) using three different values of β :
 504 0.2, 0.5, and 0.8. For comparison, we also include standard PPO without ADARFT (denoted as “w/o
 505 ADARFT”) as a baseline.

506 As shown in Figure 4, the model trained with $\beta = 0.5$ achieves the highest accuracy throughout
 507 training. This supports our theoretical motivation in Section 3.4: maximizing reward variance, which
 508 occurs when success rate ≈ 0.5 , provides the strongest learning signal. Models with $\beta = 0.2$ and
 509 $\beta = 0.8$ underperform likely due to curriculum misalignment: $\beta = 0.8$ overly focuses on easy
 510 problems, while $\beta = 0.2$ overemphasizes difficult ones, both of which limit the model’s capacity to
 511 generalize. The reward and difficulty curves align with the accuracy outcomes discussed above. The
 512 $\beta = 0.5$ configuration maintains a stable reward near 0.5, reflecting balanced difficulty exposure.
 513 In contrast, $\beta = 0.8$ results in overly high reward (i.e., easy samples), while $\beta = 0.2$ maintains a
 514 reward around 0.2 for most of training, indicating the model is repeatedly presented with overly

486 difficult problems. As expected, response length is the shortest for $\beta = 0.8$ and longest for $\beta = 0.2$,
 487 consistent with the idea that longer responses correlate with problem complexity.
 488



498 Figure 4: Ablation on β in ADARFT: we compare model accuracy, average reward, response length,
 499 and mean difficulty under $\beta = 0.2$, $\beta = 0.5$, and $\beta = 0.8$, along with standard PPO (w/o ADARFT).
 500
 501

502 6 WHEN DOES CURRICULUM LEARNING HELP?

504 Our findings show that curriculum learning provides the greatest benefits under two key conditions:
 505 (1) imbalanced training distributions, and (2) limited model capacity. In skewed distributions,
 506 particularly the skew-difficult settings, standard PPO often struggles to gain traction early in training
 507 due to insufficient reward signals. ADARFT mitigates this by initially sampling easier problems,
 508 enabling the model to bootstrap capabilities before tackling harder content. Conversely, the benefits of
 509 ADARFT are less pronounced when the model is strong enough or the data is already well-balanced.
 510 In both cases, the model is either already exposed to a representative distribution of task difficulties
 511 or finds most problems challenging enough, thus reducing the need for dynamic difficulty adjustment.
 512 In addition, we conducted further experiments detailed in Appendix A, including evaluations on
 513 datasets with more extreme difficulty distributions (A.2), difficulty estimation using an LLM-based
 514 judge (A.4), and instantiations of ADARFT with alternative RL algorithms (GRPO, REINFORCE++)
 515 (A.3). Across all these settings, ADARFT consistently demonstrates effectiveness, highlighting its
 516 robustness to diverse data distributions, compatibility with various RL algorithms, and flexibility with
 517 different difficulty metrics.

518 It is important to note that manual data curation and task scheduling carefully designed for a specific
 519 model could potentially achieve similar results by selecting a training sequence that aligns with
 520 the model’s learning capacity (Yu et al., 2025; Shen et al., 2025; Chen et al., 2025; Zeng et al.,
 521 2025). However, such methods require significant human effort, domain knowledge, and often need
 522 to be re-tuned for each new model or task. **Fixed curriculum schedules face similar limitations:**
 523 **because they rely on a predetermined ordering of difficulty, they cannot adjust when the model**
 524 **learns faster or slower than expected.** In contrast, ADARFT requires no manual data curation or
 525 model-specific preprocessing. The curriculum automatically adapts to the model’s reward signal
 526 during training, making it broadly applicable across model scales and training distributions. This
 527 automatic adaptability not only saves engineering effort but also improves scalability and robustness
 528 in real-world training pipelines. Moreover, ADARFT is particularly advantageous in fixed data
 529 settings, as it can tailor the difficulty schedule to match the capabilities of any model, whether strong
 530 or weak, without altering the dataset, providing a unified solution that generalizes across model
 531 architectures and skill levels.

532 7 CONCLUSION

534 We propose ADARFT, an adaptive curriculum learning strategy for reinforcement finetuning (RFT)
 535 that dynamically matches problem difficulty to a model’s evolving skill level. By adjusting a target
 536 difficulty based on reward feedback, ADARFT improves both sample and compute efficiency without
 537 modifying the reward function or underlying RL algorithm. Experiments across multiple data
 538 regimes and model sizes show consistent gains in convergence speed and final accuracy, especially
 539 in imbalanced training distributions. This lightweight, scalable approach highlights the value of
 curriculum-aware training for efficient and robust alignment in structured reasoning tasks.

540 REFERENCES
541

542 Arash Ahmadian, Chris Cremer, Matthias Gallé, Marzieh Fadaee, Julia Kreutzer, Olivier Pietquin,
543 Ahmet Üstün, and Sara Hooker. Back to basics: Revisiting reinforce style optimization for learning
544 from human feedback in llms, 2024. URL <https://arxiv.org/abs/2402.14740>.

545 Sanghwan Bae, Jiwoo Hong, Min Young Lee, Hanbyul Kim, JeongYeon Nam, and Donghyun
546 Kwak. Online difficulty filtering for reasoning oriented reinforcement learning. *arXiv preprint*
547 *arXiv:2504.03380*, 2025.

548

549 Yoshua Bengio, Jérôme Louradour, Ronan Collobert, and Jason Weston. Curriculum learning. In
550 *Proceedings of the 26th Annual International Conference on Machine Learning*, ICML '09, pp.
551 41–48, New York, NY, USA, 2009. Association for Computing Machinery. ISBN 9781605585161.
552 doi: 10.1145/1553374.1553380. URL <https://doi.org/10.1145/1553374.1553380>.

553 Zhipeng Chen, Yingqian Min, Beichen Zhang, Jie Chen, Jinhao Jiang, Daixuan Cheng, Wayne Xin
554 Zhao, Zheng Liu, Xu Miao, Yang Lu, Lei Fang, Zhongyuan Wang, and Ji-Rong Wen. An empirical
555 study on eliciting and improving r1-like reasoning models, 2025. URL <https://arxiv.org/abs/2503.04548>.

556

557 Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser,
558 Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John
559 Schulman. Training verifiers to solve math word problems, 2021. URL <https://arxiv.org/abs/2110.14168>.

560

561 Ganqu Cui, Lifan Yuan, Zefan Wang, Hanbin Wang, Wendi Li, Bingxiang He, Yuchen Fan, Tianyu
562 Yu, Qixin Xu, Weize Chen, Jiarui Yuan, Huayu Chen, Kaiyan Zhang, Xingtai Lv, Shuo Wang,
563 Yuan Yao, Xu Han, Hao Peng, Yu Cheng, Zhiyuan Liu, Maosong Sun, Bowen Zhou, and Ning
564 Ding. Process reinforcement through implicit rewards, 2025. URL <https://arxiv.org/abs/2502.01456>.

565

566 Wojciech Marian Czarnecki, Siddhant M. Jayakumar, Max Jaderberg, Leonard Hasenclever,
567 Yee Whye Teh, Simon Osindero, Nicolas Heess, and Razvan Pascanu. Mix&match - agent
568 curricula for reinforcement learning, 2018. URL <https://arxiv.org/abs/1806.01780>.

569

570 DeepSeek-AI, Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu,
571 Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, Xiaokang Zhang, Xingkai Yu, Yu Wu, Z. F. Wu,
572 Zhibin Gou, Zhihong Shao, Zhuoshu Li, Ziyi Gao, Aixin Liu, Bing Xue, Bingxuan Wang, Bochao
573 Wu, Bei Feng, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan,
574 Damai Dai, Deli Chen, Dongjie Ji, Erhang Li, Fangyun Lin, Fucong Dai, Fuli Luo, Guangbo Hao,
575 Guanting Chen, Guowei Li, H. Zhang, Han Bao, Hanwei Xu, Haocheng Wang, Honghui Ding,
576 Huajian Xin, Huazuo Gao, Hui Qu, Hui Li, Jianzhong Guo, Jiashi Li, Jiawei Wang, Jingchang
577 Chen, Jingyang Yuan, Junjie Qiu, Junlong Li, J. L. Cai, Jiaqi Ni, Jian Liang, Jin Chen, Kai Dong,
578 Kai Hu, Kaige Gao, Kang Guan, Kexin Huang, Kuai Yu, Lean Wang, Lecong Zhang, Liang Zhao,
579 Litong Wang, Liyue Zhang, Lei Xu, Leyi Xia, Mingchuan Zhang, Minghua Zhang, Minghui Tang,
580 Meng Li, Miaojun Wang, Mingming Li, Ning Tian, Panpan Huang, Peng Zhang, Qiancheng Wang,
581 Qinyu Chen, Qiushi Du, Ruiqi Ge, Ruisong Zhang, Ruizhe Pan, Runji Wang, R. J. Chen, R. L.
582 Jin, Ruyi Chen, Shanghao Lu, Shangyan Zhou, Shanhuan Chen, Shengfeng Ye, Shiyu Wang,
583 Shuiping Yu, Shunfeng Zhou, Shuting Pan, S. S. Li, Shuang Zhou, Shaoqing Wu, Shengfeng
584 Ye, Tao Yun, Tian Pei, Tianyu Sun, T. Wang, Wangding Zeng, Wanjia Zhao, Wen Liu, Wenfeng
585 Liang, Wenjun Gao, Wenqin Yu, Wentao Zhang, W. L. Xiao, Wei An, Xiaodong Liu, Xiaohan
586 Wang, Xiaokang Chen, Xiaotao Nie, Xin Cheng, Xin Liu, Xin Xie, Xingchao Liu, Xinyu Yang,
587 Xinyuan Li, Xuecheng Su, Xuheng Lin, X. Q. Li, Xiangyue Jin, Xiaojin Shen, Xiaosha Chen,
588 Xiaowen Sun, Xiaoxiang Wang, Xinnan Song, Xinyi Zhou, Xianzu Wang, Xinxia Shan, Y. K. Li,
589 Y. Q. Wang, Y. X. Wei, Yang Zhang, Yanhong Xu, Yao Li, Yao Zhao, Yaofeng Sun, Yaohui Wang,
590 Yi Yu, Yichao Zhang, Yifan Shi, Yiliang Xiong, Ying He, Yishi Piao, Yisong Wang, Yixuan Tan,
591 Yiyang Ma, Yiyuan Liu, Yongqiang Guo, Yuan Ou, Yuduan Wang, Yue Gong, Yuheng Zou, Yujia
592 He, Yunfan Xiong, Yuxiang Luo, Yuxiang You, Yuxuan Liu, Yuyang Zhou, Y. X. Zhu, Yanhong
593 Xu, Yanping Huang, Yaohui Li, Yi Zheng, Yuchen Zhu, Yunxian Ma, Ying Tang, Yukun Zha,
Yuting Yan, Z. Z. Ren, Zehui Ren, Zhangli Sha, Zhe Fu, Zhean Xu, Zhenda Xie, Zhengyan Zhang,
Zhewen Hao, Zhicheng Ma, Zhigang Yan, Zhiyu Wu, Zihui Gu, Zijia Zhu, Zijun Liu, Zilin Li,

594 Ziwei Xie, Ziyang Song, Zizheng Pan, Zhen Huang, Zhipeng Xu, Zhongyu Zhang, and Zhen
 595 Zhang. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning, 2025.
 596 URL <https://arxiv.org/abs/2501.12948>.

597
 598 Hanze Dong, Wei Xiong, Deepanshu Goyal, Yihan Zhang, Winnie Chow, Rui Pan, Shizhe Diao,
 599 Jipeng Zhang, Kashun Shum, and Tong Zhang. Raft: Reward ranked finetuning for generative
 600 foundation model alignment, 2023. URL <https://arxiv.org/abs/2304.06767>.

601
 602 Carlos Florensa, David Held, Xinyang Geng, and Pieter Abbeel. Automatic goal generation for
 603 reinforcement learning agents, 2018. URL <https://arxiv.org/abs/1705.06366>.

604
 605 Bofei Gao, Feifan Song, Zhe Yang, Zefan Cai, Yibo Miao, Qingxiu Dong, Lei Li, Chenghao Ma,
 606 Liang Chen, Runxin Xu, Zhengyang Tang, Benyou Wang, Daoguang Zan, Shanghaoran Quan,
 607 Ge Zhang, Lei Sha, Yichang Zhang, Xuancheng Ren, Tianyu Liu, and Baobao Chang. Omnimath:
 608 A universal olympiad level mathematic benchmark for large language models, 2024. URL
<https://arxiv.org/abs/2410.07985>.

609
 610 Dongyoung Go, Tomasz Korbak, Germán Kruszewski, Jos Rozen, Nahyeon Ryu, and Marc Dymet-
 611 man. Aligning language models with preferences through f-divergence minimization. *arXiv*
 612 preprint *arXiv:2302.08215*, 2023.

613 Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad
 614 Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, Amy Yang, Angela Fan,
 615 Anirudh Goyal, Anthony Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev,
 616 Arthur Hinsvark, Arun Rao, Aston Zhang, Aurelien Rodriguez, Austen Gregerson, Ava Spataru,
 617 Baptiste Roziere, Bethany Biron, Binh Tang, Bobbie Chern, Charlotte Caucheteux, Chaya Nayak,
 618 Chloe Bi, Chris Marra, Chris McConnell, Christian Keller, Christophe Touret, Chunyang Wu,
 619 Corinne Wong, Cristian Canton Ferrer, Cyrus Nikolaidis, Damien Allonsius, Daniel Song, Danielle
 620 Pintz, Danny Livshits, Danny Wyatt, David Esiobu, Dhruv Choudhary, Dhruv Mahajan, Diego
 621 Garcia-Olano, Diego Perino, Dieuwke Hupkes, Egor Lakomkin, Ehab AlBadawy, Elina Lobanova,
 622 Emily Dinan, Eric Michael Smith, Filip Radenovic, Francisco Guzmán, Frank Zhang, Gabriel
 623 Synnaeve, Gabrielle Lee, Georgia Lewis Anderson, Govind Thattai, Graeme Nail, Gregoire Mialon,
 624 Guan Pang, Guillem Cucurell, Hailey Nguyen, Hannah Korevaar, Hu Xu, Hugo Touvron, Iliyan
 625 Zarov, Imanol Arrieta Ibarra, Isabel Kloumann, Ishan Misra, Ivan Evtimov, Jack Zhang, Jade Copet,
 626 Jaewon Lee, Jan Geffert, Jana Vrane, Jason Park, Jay Mahadeokar, Jeet Shah, Jelmer van der Linde,
 627 Jennifer Billock, Jenny Hong, Jenya Lee, Jeremy Fu, Jianfeng Chi, Jianyu Huang, Jiawen Liu, Jie
 628 Wang, Jiecao Yu, Joanna Bitton, Joe Spisak, Jongsoo Park, Joseph Rocca, Joshua Johnstun, Joshua
 629 Saxe, Junteng Jia, Kalyan Vasudevan Alwala, Karthik Prasad, Kartikeya Upasani, Kate Plawiak,
 630 Ke Li, Kenneth Heafield, Kevin Stone, Khalid El-Arini, Krithika Iyer, Kshitiz Malik, Kuenley
 631 Chiu, Kunal Bhalla, Kushal Lakhota, Lauren Rantala-Yearly, Laurens van der Maaten, Lawrence
 632 Chen, Liang Tan, Liz Jenkins, Louis Martin, Lovish Madaan, Lubo Malo, Lukas Blecher, Lukas
 633 Landzaat, Luke de Oliveira, Madeline Muzzi, Mahesh Pasupuleti, Mannat Singh, Manohar Paluri,
 634 Marcin Kardas, Maria Tsimpoukelli, Mathew Oldham, Mathieu Rita, Maya Pavlova, Melanie
 635 Kambadur, Mike Lewis, Min Si, Mitesh Kumar Singh, Mona Hassan, Naman Goyal, Narjes
 636 Torabi, Nikolay Bashlykov, Nikolay Bogoychev, Niladri Chatterji, Ning Zhang, Olivier Duchenne,
 637 Onur Çelebi, Patrick Alrassy, Pengchuan Zhang, Pengwei Li, Petar Vasic, Peter Weng, Prajwal
 638 Bhargava, Pratik Dubal, Praveen Krishnan, Punit Singh Koura, Puxin Xu, Qing He, Qingxiao Dong,
 639 Ragavan Srinivasan, Raj Ganapathy, Ramon Calderer, Ricardo Silveira Cabral, Robert Stojnic,
 640 Roberta Raileanu, Rohan Maheswari, Rohit Girdhar, Rohit Patel, Romain Sauvestre, Ronnie
 641 Polidoro, Roshan Sumbaly, Ross Taylor, Ruan Silva, Rui Hou, Rui Wang, Saghar Hosseini, Sahana
 642 Chennabasappa, Sanjay Singh, Sean Bell, Seohyun Sonia Kim, Sergey Edunov, Shaoliang Nie,
 643 Sharan Narang, Sharath Raparth, Sheng Shen, Shengye Wan, Shruti Bhosale, Shun Zhang, Simon
 644 Vandenhende, Soumya Batra, Spencer Whitman, Sten Sootla, Stephane Collot, Suchin Gururangan,
 645 Sydney Borodinsky, Tamar Herman, Tara Fowler, Tarek Sheasha, Thomas Georgiou, Thomas
 646 Scialom, Tobias Speckbacher, Todor Mihaylov, Tong Xiao, Ujjwal Karn, Vedanuj Goswami,
 647 Vibhor Gupta, Vignesh Ramanathan, Viktor Kerkez, Vincent Gonguet, Virginie Do, Vish Vogeti,
 648 Vitor Albiero, Vladan Petrovic, Weiwei Chu, Wenhan Xiong, Wenyan Fu, Whitney Meers, Xavier
 649 Martinet, Xiaodong Wang, Xiaofang Wang, Xiaoqing Ellen Tan, Xide Xia, Xinfeng Xie, Xuchao
 650 Jia, Xuewei Wang, Yaelle Goldschlag, Yashesh Gaur, Yasmine Babaei, Yi Wen, Yiwen Song,
 651 Yuchen Zhang, Yue Li, Yuning Mao, Zacharie Delpierre Coudert, Zheng Yan, Zhengxing Chen, Zoe

648 Papakipos, Aaditya Singh, Aayushi Srivastava, Abha Jain, Adam Kelsey, Adam Shajnfeld, Adithya
 649 Gangidi, Adolfo Victoria, Ahuva Goldstand, Ajay Menon, Ajay Sharma, Alex Boesenber, Alexei
 650 Baevski, Allie Feinstein, Amanda Kallet, Amit Sangani, Amos Teo, Anam Yunus, Andrei Lupu,
 651 Andres Alvarado, Andrew Caples, Andrew Gu, Andrew Ho, Andrew Poulton, Andrew Ryan, Ankit
 652 Ramchandani, Annie Dong, Annie Franco, Anuj Goyal, Aparajita Saraf, Arkabandhu Chowdhury,
 653 Ashley Gabriel, Ashwin Bharambe, Assaf Eisenman, Azadeh Yazdan, Beau James, Ben Maurer,
 654 Benjamin Leonhardi, Bernie Huang, Beth Loyd, Beto De Paola, Bhargavi Paranjape, Bing Liu,
 655 Bo Wu, Boyu Ni, Braden Hancock, Bram Wasti, Brandon Spence, Brani Stojkovic, Brian Gamido,
 656 Britt Montalvo, Carl Parker, Carly Burton, Catalina Mejia, Ce Liu, Changhan Wang, Changkyu
 657 Kim, Chao Zhou, Chester Hu, Ching-Hsiang Chu, Chris Cai, Chris Tindal, Christoph Feichtenhofer,
 658 Cynthia Gao, Damon Civin, Dana Beaty, Daniel Kreymer, Daniel Li, David Adkins, David Xu,
 659 Davide Testuggine, Delia David, Devi Parikh, Diana Liskovich, Didem Foss, Dingkang Wang, Duc
 660 Le, Dustin Holland, Edward Dowling, Eissa Jamil, Elaine Montgomery, Eleonora Presani, Emily
 661 Hahn, Emily Wood, Eric-Tuan Le, Erik Brinkman, Esteban Arcaute, Evan Dunbar, Evan Smothers,
 662 Fei Sun, Felix Kreuk, Feng Tian, Filippos Kokkinos, Firat Ozgenel, Francesco Caggioni, Frank
 663 Kanayet, Frank Seide, Gabriela Medina Florez, Gabriella Schwarz, Gada Badeer, Georgia Swee,
 664 Gil Halpern, Grant Herman, Grigory Sizov, Guangyi, Zhang, Guna Lakshminarayanan, Hakan Inan,
 665 Hamid Shojanazeri, Han Zou, Hannah Wang, Hanwen Zha, Haroun Habeeb, Harrison Rudolph,
 666 Helen Suk, Henry Aspren, Hunter Goldman, Hongyuan Zhan, Ibrahim Damlaj, Igor Molybog,
 667 Igor Tufanov, Ilias Leontiadis, Irina-Elena Veliche, Itai Gat, Jake Weissman, James Geboski, James
 668 Kohli, Janice Lam, Japhet Asher, Jean-Baptiste Gaya, Jeff Marcus, Jeff Tang, Jennifer Chan, Jenny
 669 Zhen, Jeremy Reizenstein, Jeremy Teboul, Jessica Zhong, Jian Jin, Jingyi Yang, Joe Cummings,
 670 Jon Carvill, Jon Shepard, Jonathan McPhie, Jonathan Torres, Josh Ginsburg, Junjie Wang, Kai
 671 Wu, Kam Hou U, Karan Saxena, Kartikay Khandelwal, Katayoun Zand, Kathy Matosich, Kaushik
 672 Veeraghavan, Kelly Michelena, Keqian Li, Kiran Jagadeesh, Kun Huang, Kunal Chawla, Kyle
 673 Huang, Lailin Chen, Lakshya Garg, Lavender A, Leandro Silva, Lee Bell, Lei Zhang, Liangpeng
 674 Guo, Licheng Yu, Liron Moshkovich, Luca Wehrstedt, Madian Khabsa, Manav Avalani, Manish
 675 Bhatt, Martynas Mankus, Matan Hasson, Matthew Lennie, Matthias Reso, Maxim Groshev, Maxim
 676 Naumov, Maya Lathi, Meghan Keneally, Miao Liu, Michael L. Seltzer, Michal Valko, Michelle
 677 Restrepo, Mihir Patel, Mik Vyatskov, Mikayel Samvelyan, Mike Clark, Mike Macey, Mike Wang,
 678 Miquel Jubert Hermoso, Mo Metanat, Mohammad Rastegari, Munish Bansal, Nandhini Santhanam,
 679 Natascha Parks, Natasha White, Navyata Bawa, Nayan Singhal, Nick Egebo, Nicolas Usunier,
 680 Nikhil Mehta, Nikolay Pavlovich Laptev, Ning Dong, Norman Cheng, Oleg Chernoguz, Olivia
 681 Hart, Omkar Salpekar, Ozlem Kalinli, Parkin Kent, Parth Parekh, Paul Saab, Pavan Balaji, Pedro
 682 Rittner, Philip Bontrager, Pierre Roux, Piotr Dollar, Polina Zvyagina, Prashant Ratanchandani,
 683 Pritish Yuvraj, Qian Liang, Rachad Alao, Rachel Rodriguez, Rafi Ayub, Raghotham Murthy,
 684 Raghu Nayani, Rahul Mitra, Rangaprabhu Parthasarathy, Raymond Li, Rebekkah Hogan, Robin
 685 Battey, Rocky Wang, Russ Howes, Ruty Rinott, Sachin Mehta, Sachin Siby, Sai Jayesh Bondu,
 686 Samyak Datta, Sara Chugh, Sara Hunt, Sargun Dhillon, Sasha Sidorov, Satadru Pan, Saurabh
 687 Mahajan, Saurabh Verma, Seiji Yamamoto, Sharadh Ramaswamy, Shaun Lindsay, Shaun Lindsay,
 688 Sheng Feng, Shenghao Lin, Shengxin Cindy Zha, Shishir Patil, Shiva Shankar, Shuqiang Zhang,
 689 Shuqiang Zhang, Sinong Wang, Sneha Agarwal, Soji Sajuyigbe, Soumith Chintala, Stephanie
 690 Max, Stephen Chen, Steve Kehoe, Steve Satterfield, Sudarshan Govindaprasad, Sumit Gupta,
 691 Summer Deng, Sungmin Cho, Sunny Virk, Suraj Subramanian, Sy Choudhury, Sydney Goldman,
 692 Tal Remez, Tamar Glaser, Tamara Best, Thilo Koehler, Thomas Robinson, Tianhe Li, Tianjun
 693 Zhang, Tim Matthews, Timothy Chou, Tzook Shaked, Varun Vontimitta, Victoria Ajayi, Victoria
 694 Montanez, Vijai Mohan, Vinay Satish Kumar, Vishal Mangla, Vlad Ionescu, Vlad Poenaru,
 695 Vlad Tiberiu Mihailescu, Vladimir Ivanov, Wei Li, Wencheng Wang, Wenwen Jiang, Wes Bouaziz,
 696 Will Constable, Xiaocheng Tang, Xiaoqian Wu, Xiaolan Wang, Xilun Wu, Xinbo Gao, Yaniv
 697 Kleinman, Yanjun Chen, Ye Hu, Ye Jia, Ye Qi, Yenda Li, Yilin Zhang, Ying Zhang, Yossi Adi,
 698 Youngjin Nam, Yu, Wang, Yu Zhao, Yuchen Hao, Yundi Qian, Yunlu Li, Yuzi He, Zach Rait,
 699 Zachary DeVito, Zef Rosnbrick, Zhaoduo Wen, Zhenyu Yang, Zhiwei Zhao, and Zhiyu Ma. The
 700 llama 3 herd of models, 2024. URL <https://arxiv.org/abs/2407.21783>.
 701 Tuomas Haarnoja, Haoran Tang, Pieter Abbeel, and Sergey Levine. Reinforcement learning with
 702 deep energy-based policies. In *International conference on machine learning*, pp. 1352–1361.
 703 PMLR, 2017.
 704 Chaoqun He, Renjie Luo, Yuzhuo Bai, Shengding Hu, Zhen Leng Thai, Junhao Shen, Jinyi Hu,
 705 Xu Han, Yujie Huang, Yuxiang Zhang, Jie Liu, Lei Qi, Zhiyuan Liu, and Maosong Sun. Olympiad-

702 bench: A challenging benchmark for promoting agi with olympiad-level bilingual multimodal
 703 scientific problems, 2024. URL <https://arxiv.org/abs/2402.14008>.

704

705 Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song,
 706 and Jacob Steinhardt. Measuring mathematical problem solving with the MATH dataset. In
 707 *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track*
 708 (*Round 2*), 2021. URL <https://openreview.net/forum?id=7Bywt2mQsCe>.

709 Jian Hu. Reinforce++: A simple and efficient approach for aligning large language models, 2025.
 710 URL <https://arxiv.org/abs/2501.03262>.

711

712 Jingcheng Hu, Yinmin Zhang, Qi Han, Dixin Jiang, Xiangyu Zhang, and Heung-Yeung Shum.
 713 Open-reasoner-zero: An open source approach to scaling up reinforcement learning on the base
 714 model, 2025. URL <https://arxiv.org/abs/2503.24290>.

715 Allan Jabri, Kyle Hsu, Ben Eysenbach, Abhishek Gupta, Sergey Levine, and Chelsea Finn. Unsu-
 716 pervised curricula for visual meta-reinforcement learning, 2019. URL <https://arxiv.org/abs/1912.04226>.

717

718 Niels Justesen, Ruben Rodriguez Torrado, Philip Bontrager, Ahmed Khalifa, Julian Togelius, and
 719 Sebastian Risi. Illuminating generalization in deep reinforcement learning through procedural
 720 level generation, 2018. URL <https://arxiv.org/abs/1806.10729>.

721

722 Amirhossein Kazemnejad, Milad Aghajohari, Eva Portelance, Alessandro Sordoni, Siva Reddy,
 723 Aaron Courville, and Nicolas Le Roux. Vineppo: Unlocking rl potential for llm reasoning through
 724 refined credit assignment, 2024. URL <https://arxiv.org/abs/2410.01679>.

725

726 Tomasz Korbak, Hady Elsahar, Germán Kruszewski, and Marc Dymetman. On reinforcement
 727 learning and distribution matching for fine-tuning language models with no catastrophic forgetting.
 728 *Advances in Neural Information Processing Systems*, 35:16203–16220, 2022.

729

730 Aitor Lewkowycz, Anders Andreassen, David Dohan, Ethan Dyer, Henryk Michalewski, Vinay
 731 Ramasesh, Ambrose Sloane, Cem Anil, Imanol Schlag, Theo Gutman-Solo, Yuhuai Wu, Behnam
 732 Neyshabur, Guy Gur-Ari, and Vedant Misra. Solving quantitative reasoning problems with language
 733 models, 2022. URL <https://arxiv.org/abs/2206.14858>.

734

735 Xuefeng Li, Haoyang Zou, and Pengfei Liu. Limr: Less is more for rl scaling, 2025. URL
 736 <https://arxiv.org/abs/2502.11886>.

737

738 Ziniu Li, Tian Xu, Yushun Zhang, Zhihang Lin, Yang Yu, Ruoyu Sun, and Zhi-Quan Luo. Remax: A
 739 simple, effective, and efficient reinforcement learning method for aligning large language models,
 740 2024. URL <https://arxiv.org/abs/2310.10505>.

741

742 Hunter Lightman, Vineet Kosaraju, Yura Burda, Harri Edwards, Bowen Baker, Teddy Lee, Jan
 743 Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. Let's verify step by step, 2023. URL
 744 <https://arxiv.org/abs/2305.20050>.

745

746 Michael Luo, Sijun Tan, Justin Wong, Xiaoxiang Shi, William Y. Tang, Manan Roongta, Colin Cai,
 747 Jeffrey Luo, Tianjun Zhang, Li Erran Li, Raluca Ada Popa, and Ion Stoica. Deepscaler: Surpassing
 748 o1-preview with a 1.5b model by scaling rl, 2025. Notion Blog.

749

750 Tambet Matiisen, Avital Oliver, Taco Cohen, and John Schulman. Teacher-student curriculum
 751 learning, 2017. URL <https://arxiv.org/abs/1707.00183>.

752

753 Niklas Muennighoff, Zitong Yang, Weijia Shi, Xiang Lisa Li, Li Fei-Fei, Hannaneh Hajishirzi, Luke
 754 Zettlemoyer, Percy Liang, Emmanuel Candès, and Tatsunori Hashimoto. s1: Simple test-time
 755 scaling, 2025. URL <https://arxiv.org/abs/2501.19393>.

756

757 OpenAI, :, Aaron Hurst, Adam Lerer, Adam P. Goucher, Adam Perelman, Aditya Ramesh, Aidan
 758 Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, Aleksander Mądry, Alex Baker-
 759 Whitcomb, Alex Beutel, Alex Borzunov, Alex Carney, Alex Chow, Alex Kirillov, Alex Nichol, Alex
 760 Paino, Alex Renzin, Alex Tachard Passos, Alexander Kirillov, Alexi Christakis, Alexis Conneau,
 761 Ali Kamali, Allan Jabri, Allison Moyer, Allison Tam, Amadou Crookes, Amin Tootoochian,

756 Amin Tootoonchian, Ananya Kumar, Andrea Vallone, Andrej Karpathy, Andrew Braunstein,
 757 Andrew Cann, Andrew Codispoti, Andrew Galu, Andrew Kondrich, Andrew Tulloch, Andrey
 758 Mishchenko, Angela Baek, Angela Jiang, Antoine Pelisse, Antonia Woodford, Anuj Gosalia,
 759 Arka Dhar, Ashley Pantuliano, Avi Nayak, Avital Oliver, Barret Zoph, Behrooz Ghorbani, Ben
 760 Leimberger, Ben Rossen, Ben Sokolowsky, Ben Wang, Benjamin Zweig, Beth Hoover, Blake
 761 Samic, Bob McGrew, Bobby Spero, Bogo Giertler, Bowen Cheng, Brad Lightcap, Brandon
 762 Walkin, Brendan Quinn, Brian Guaraci, Brian Hsu, Bright Kellogg, Brydon Eastman, Camillo
 763 Lugaressi, Carroll Wainwright, Cary Bassin, Cary Hudson, Casey Chu, Chad Nelson, Chak Li,
 764 Chan Jun Shern, Channing Conger, Charlotte Barette, Chelsea Voss, Chen Ding, Cheng Lu,
 765 Chong Zhang, Chris Beaumont, Chris Hallacy, Chris Koch, Christian Gibson, Christina Kim,
 766 Christine Choi, Christine McLeavey, Christopher Hesse, Claudia Fischer, Clemens Winter, Coley
 767 Czarnecki, Colin Jarvis, Colin Wei, Constantin Koumouzelis, Dane Sherburn, Daniel Kappler,
 768 Daniel Levin, Daniel Levy, David Carr, David Farhi, David Mely, David Robinson, David Sasaki,
 769 Denny Jin, Dev Valladares, Dimitris Tsipras, Doug Li, Duc Phong Nguyen, Duncan Findlay,
 770 Edede Oiwoh, Edmund Wong, Ehsan Asdar, Elizabeth Proehl, Elizabeth Yang, Eric Antonow, Eric
 771 Kramer, Eric Peterson, Eric Sigler, Eric Wallace, Eugene Brevdo, Evan Mays, Farzad Khorasani,
 772 Felipe Petroski Such, Filippo Rasoppi, Francis Zhang, Fred von Lohmann, Freddie Sulit, Gabriel Goh,
 773 Gene Oden, Geoff Salmon, Giulio Starace, Greg Brockman, Hadi Salman, Haiming Bao, Haitang
 774 Hu, Hannah Wong, Haoyu Wang, Heather Schmidt, Heather Whitney, Heewoo Jun, Hendrik
 775 Kirchner, Henrique Ponde de Oliveira Pinto, Hongyu Ren, Huiwen Chang, Hyung Won Chung,
 776 Ian Kivlichan, Ian O'Connell, Ian O'Connell, Ian Osband, Ian Silber, Ian Sohl, Ibrahim Okuyucu,
 777 Ikai Lan, Ilya Kostrikov, Ilya Sutskever, Ingmar Kanitscheider, Ishaan Gulrajani, Jacob Coxon,
 778 Jacob Menick, Jakub Pachocki, James Aung, James Betker, James Crooks, James Lennon, Jamie
 779 Kiros, Jan Leike, Jane Park, Jason Kwon, Jason Phang, Jason Teplitz, Jason Wei, Jason Wolfe,
 780 Jay Chen, Jeff Harris, Jenia Varavva, Jessica Gan Lee, Jessica Shieh, Ji Lin, Jiahui Yu, Jiayi
 781 Weng, Jie Tang, Jieqi Yu, Joanne Jang, Joaquin Quinonero Candela, Joe Beutler, Joe Landers,
 782 Joel Parish, Johannes Heidecke, John Schulman, Jonathan Lachman, Jonathan McKay, Jonathan
 783 Uesato, Jonathan Ward, Jong Wook Kim, Joost Huizinga, Jordan Sitkin, Jos Kraaijeveld, Josh
 784 Gross, Josh Kaplan, Josh Snyder, Joshua Achiam, Joy Jiao, Joyce Lee, Juntang Zhuang, Justyn
 785 Harriman, Kai Fricke, Kai Hayashi, Karan Singh, Katy Shi, Kavin Karthik, Kayla Wood, Kendra
 786 Rimbach, Kenny Hsu, Kenny Nguyen, Keren Gu-Lemberg, Kevin Button, Kevin Liu, Kiel Howe,
 787 Krithika Muthukumar, Kyle Luther, Lama Ahmad, Larry Kai, Lauren Itow, Lauren Workman,
 788 Leher Pathak, Leo Chen, Li Jing, Lia Guy, Liam Fedus, Liang Zhou, Lien Mamitsuka, Lilian Weng,
 789 Lindsay McCallum, Lindsey Held, Long Ouyang, Louis Feuvrier, Lu Zhang, Lukas Kondraciuk,
 790 Lukasz Kaiser, Luke Hewitt, Luke Metz, Lyric Doshi, Mada Aflak, Maddie Simens, Madelaine
 791 Boyd, Madeleine Thompson, Marat Dukhan, Mark Chen, Mark Gray, Mark Hudnall, Marvin
 792 Zhang, Marwan Aljubeh, Mateusz Litwin, Matthew Zeng, Max Johnson, Maya Shetty, Mayank
 793 Gupta, Meghan Shah, Mehmet Yatbaz, Meng Jia Yang, Mengchao Zhong, Mia Glaese, Mianna
 794 Chen, Michael Janner, Michael Lampe, Michael Petrov, Michael Wu, Michele Wang, Michelle
 795 Fradin, Michelle Pokrass, Miguel Castro, Miguel Oom Temudo de Castro, Mikhail Pavlov, Miles
 796 Brundage, Miles Wang, Minal Khan, Mira Murati, Mo Bavarian, Molly Lin, Murat Yesildal, Nacho
 797 Soto, Natalia Gimelshein, Natalie Cone, Natalie Staudacher, Natalie Summers, Natan LaFontaine,
 798 Neil Chowdhury, Nick Ryder, Nick Stathas, Nick Turley, Nik Tezak, Niko Felix, Nithanth Kudige,
 799 Nitish Keskar, Noah Deutsch, Noel Bundick, Nora Puckett, Ofir Nachum, Ola Okelola, Oleg Boiko,
 800 Oleg Murk, Oliver Jaffe, Olivia Watkins, Olivier Godement, Owen Campbell-Moore, Patrick
 801 Chao, Paul McMillan, Pavel Belov, Peng Su, Peter Bak, Peter Bakkum, Peter Deng, Peter Dolan,
 802 Peter Hoeschele, Peter Welinder, Phil Tillet, Philip Pronin, Philippe Tillet, Prafulla Dhariwal,
 803 Qiming Yuan, Rachel Dias, Rachel Lim, Rahul Arora, Rajan Troll, Randall Lin, Rapha Gontijo
 804 Lopes, Raul Puri, Reah Miyara, Reimar Leike, Renaud Gaubert, Reza Zamani, Ricky Wang, Rob
 805 Donnelly, Rob Honsby, Rocky Smith, Rohan Sahai, Rohit Ramchandani, Romain Huet, Rory
 806 Carmichael, Rowan Zellers, Roy Chen, Ruby Chen, Ruslan Nigmatullin, Ryan Cheu, Saachi
 807 Jain, Sam Altman, Sam Schoenholz, Sam Toizer, Samuel Miserendino, Sandhini Agarwal, Sara
 808 Culver, Scott Ethersmith, Scott Gray, Sean Grove, Sean Metzger, Shamez Hermani, Shantanu
 809 Jain, Shengjia Zhao, Sherwin Wu, Shino Jomoto, Shirong Wu, Shuaiqi Xia, Sonia Phene, Spencer
 Papay, Srinivas Narayanan, Steve Coffey, Steve Lee, Stewart Hall, Suchir Balaji, Tal Broda, Tal
 Stramer, Tao Xu, Tarun Gogineni, Taya Christianson, Ted Sanders, Tejal Patwardhan, Thomas
 Cunningham, Thomas Degry, Thomas Dimson, Thomas Raoux, Thomas Shadwell, Tianhao
 Zheng, Todd Underwood, Todor Markov, Toki Sherbakov, Tom Rubin, Tom Stasi, Tomer Kaftan,
 Tristan Heywood, Troy Peterson, Tyce Walters, Tyna Eloundou, Valerie Qi, Veit Moeller, Vinnie

810 Monaco, Vishal Kuo, Vlad Fomenko, Wayne Chang, Weiyi Zheng, Wenda Zhou, Wesam Manassra,
 811 Will Sheu, Wojciech Zaremba, Yash Patil, Yilei Qian, Yongjik Kim, Youlong Cheng, Yu Zhang,
 812 Yuchen He, Yuchen Zhang, Yujia Jin, Yunxing Dai, and Yury Malkov. Gpt-4o system card, 2024a.
 813 URL <https://arxiv.org/abs/2410.21276>.

814 OpenAI, :, Aaron Jaech, Adam Kalai, Adam Lerer, Adam Richardson, Ahmed El-Kishky, Aiden
 815 Low, Alec Helyar, Aleksander Madry, Alex Beutel, Alex Carney, Alex Iftimie, Alex Karpenko,
 816 Alex Tachard Passos, Alexander Neitz, Alexander Prokofiev, Alexander Wei, Allison Tam, Ally
 817 Bennett, Ananya Kumar, Andre Saraiva, Andrea Vallone, Andrew Duberstein, Andrew Kondrich,
 818 Andrey Mishchenko, Andy Applebaum, Angela Jiang, Ashvin Nair, Barret Zoph, Behrooz Ghor-
 819 bani, Ben Rossen, Benjamin Sokolowsky, Boaz Barak, Bob McGrew, Borys Minaiev, Botao
 820 Hao, Bowen Baker, Brandon Houghton, Brandon McKinzie, Brydon Eastman, Camillo Lugaresi,
 821 Cary Bassin, Cary Hudson, Chak Ming Li, Charles de Bourcy, Chelsea Voss, Chen Shen, Chong
 822 Zhang, Chris Koch, Chris Orsinger, Christopher Hesse, Claudia Fischer, Clive Chan, Dan Roberts,
 823 Daniel Kappler, Daniel Levy, Daniel Selsam, David Dohan, David Farhi, David Mely, David
 824 Robinson, Dimitris Tsipras, Doug Li, Dragos Oprica, Eben Freeman, Eddie Zhang, Edmund Wong,
 825 Elizabeth Proehl, Enoch Cheung, Eric Mitchell, Eric Wallace, Erik Ritter, Evan Mays, Fan Wang,
 826 Felipe Petroski Such, Filippo Raso, Florencia Leoni, Foivos Tsimpourlas, Francis Song, Fred
 827 von Lohmann, Freddie Sulit, Geoff Salmon, Giambattista Parascandolo, Gildas Chabot, Grace
 828 Zhao, Greg Brockman, Guillaume Leclerc, Hadi Salman, Haiming Bao, Hao Sheng, Hart Andrin,
 829 Hessam Bagherinezhad, Hongyu Ren, Hunter Lightman, Hyung Won Chung, Ian Kivlichan, Ian
 830 O'Connell, Ian Osband, Ignasi Clavera Gilaberte, Ilge Akkaya, Ilya Kostrikov, Ilya Sutskever,
 831 Irina Kofman, Jakub Pachocki, James Lennon, Jason Wei, Jean Harb, Jerry Twore, Jiacheng Feng,
 832 Jiahui Yu, Jiayi Weng, Jie Tang, Jieqi Yu, Joaquin Quiñonero Candela, Joe Palermo, Joel Parish,
 833 Johannes Heidecke, John Hallman, John Rizzo, Jonathan Gordon, Jonathan Uesato, Jonathan
 834 Ward, Joost Huizinga, Julie Wang, Kai Chen, Kai Xiao, Karan Singhal, Karina Nguyen, Karl
 835 Cobbe, Katy Shi, Kayla Wood, Kendra Rimbach, Keren Gu-Lemberg, Kevin Liu, Kevin Lu, Kevin
 836 Stone, Kevin Yu, Lama Ahmad, Lauren Yang, Leo Liu, Leon Maksin, Leyton Ho, Liam Fedus,
 837 Lilian Weng, Linden Li, Lindsay McCallum, Lindsey Held, Lorenz Kuhn, Lukas Kondraciuk,
 838 Lukasz Kaiser, Luke Metz, Madelaine Boyd, Maja Trebacz, Manas Joglekar, Mark Chen, Marko
 839 Tintor, Mason Meyer, Matt Jones, Matt Kaufer, Max Schwarzer, Meghan Shah, Mehmet Yatbaz,
 840 Melody Y. Guan, Mengyuan Xu, Mengyuan Yan, Mia Glaese, Mianna Chen, Michael Lampe,
 841 Michael Malek, Michele Wang, Michelle Fradin, Mike McClay, Mikhail Pavlov, Miles Wang,
 842 Mingxuan Wang, Mira Murati, Mo Bavarian, Mostafa Rohaninejad, Nat McAleese, Neil Chowd-
 843 hury, Neil Chowdhury, Nick Ryder, Nikolas Tezak, Noam Brown, Ofir Nachum, Oleg Boiko, Oleg
 844 Murk, Olivia Watkins, Patrick Chao, Paul Ashbourne, Pavel Izmailov, Peter Zhokhov, Rachel Dias,
 845 Rahul Arora, Randall Lin, Rapha Gontijo Lopes, Raz Gaon, Reah Miyara, Reimar Leike, Renny
 846 Hwang, Rhythm Garg, Robin Brown, Roshan James, Rui Shu, Ryan Cheu, Ryan Greene, Saachi
 847 Jain, Sam Altman, Sam Toizer, Sam Toyer, Samuel Miserendino, Sandhini Agarwal, Santiago
 848 Hernandez, Sasha Baker, Scott McKinney, Scottie Yan, Shengjia Zhao, Shengli Hu, Shibani
 849 Santurkar, Shraman Ray Chaudhuri, Shuyuan Zhang, Siyuan Fu, Spencer Papay, Steph Lin, Suchir
 850 Balaji, Suvansh Sanjeev, Szymon Sidor, Tal Broda, Aidan Clark, Tao Wang, Taylor Gordon, Ted
 851 Sanders, Tejal Patwardhan, Thibault Sottiaux, Thomas Degry, Thomas Dimson, Tianhao Zheng,
 852 Timur Garipov, Tom Stasi, Trapit Bansal, Trevor Creech, Troy Peterson, Tyna Eloundou, Valerie
 853 Qi, Vineet Kosaraju, Vinnie Monaco, Vitchyr Pong, Vlad Fomenko, Weiyi Zheng, Wenda Zhou,
 854 Wes McCabe, Wojciech Zaremba, Yann Dubois, Yinghai Lu, Yining Chen, Young Cha, Yu Bai,
 855 Yuchen He, Yuchen Zhang, Yunyun Wang, Zheng Shao, and Zhuohan Li. Openai o1 system card,
 856 2024b. URL <https://arxiv.org/abs/2412.16720>.

857 Shubham Parashar, Shurui Gui, Xiner Li, Hongyi Ling, Sushil Vemuri, Blake Olson, Eric Li,
 858 Yu Zhang, James Caverlee, Dileep Kalathil, and Shuiwang Ji. Curriculum reinforcement learning
 859 from easy to hard tasks improves llm reasoning, 2025. URL <https://arxiv.org/abs/2506.06632>.

860 Rémy Portelas, Cédric Colas, Katja Hofmann, and Pierre-Yves Oudeyer. Teacher algorithms
 861 for curriculum learning of deep rl in continuously parameterized environments, 2019. URL
 862 <https://arxiv.org/abs/1910.07224>.

863 Qwen, :, An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan
 864 Li, Dayiheng Liu, Fei Huang, Haoran Wei, Huan Lin, Jian Yang, Jianhong Tu, Jianwei Zhang,

864 Jianxin Yang, Jiaxi Yang, Jingren Zhou, Junyang Lin, Kai Dang, Keming Lu, Keqin Bao, Kexin
 865 Yang, Le Yu, Mei Li, Mingfeng Xue, Pei Zhang, Qin Zhu, Rui Men, Runji Lin, Tianhao Li, Tianyi
 866 Tang, Tingyu Xia, Xingzhang Ren, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yu Wan,
 867 Yuqiong Liu, Zeyu Cui, Zhenru Zhang, and Zihan Qiu. Qwen2.5 technical report, 2025. URL
 868 <https://arxiv.org/abs/2412.15115>.

869 Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea
 870 Finn. Direct preference optimization: Your language model is secretly a reward model. *Advances*
 871 in *Neural Information Processing Systems*, 36:53728–53741, 2023.

872 Rafael Rafailov, Joey Hejna, Ryan Park, and Chelsea Finn. From r to q^* : Your language model is
 873 secretly a q -function. *arXiv preprint arXiv:2404.12358*, 2024.

874 Pierre Harvey Richemond, Yunhao Tang, Daniel Guo, Daniele Calandriello, Mohammad Gheshlaghi
 875 Azar, Rafael Rafailov, Bernardo Avila Pires, Eugene Tarassov, Lucas Spangher, Will Ellsworth,
 876 et al. Offline regularised reinforcement learning for large language models alignment. *arXiv*
 877 *preprint arXiv:2405.19107*, 2024.

878 Andrei A. Rusu, Neil C. Rabinowitz, Guillaume Desjardins, Hubert Soyer, James Kirkpatrick, Koray
 879 Kavukcuoglu, Razvan Pascanu, and Raia Hadsell. Progressive neural networks, 2022. URL
 880 <https://arxiv.org/abs/1606.04671>.

881 John Schulman, Xi Chen, and Pieter Abbeel. Equivalence between policy gradients and soft q-learning.
 882 *arXiv preprint arXiv:1704.06440*, 2017a.

883 John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy
 884 optimization algorithms, 2017b. URL <https://arxiv.org/abs/1707.06347>.

885 Wei Shen, Guanlin Liu, Zheng Wu, Ruofei Zhu, Qingping Yang, Chao Xin, Yu Yue, and Lin Yan.
 886 Exploring data scaling trends and effects in reinforcement learning from human feedback, 2025.
 887 URL <https://arxiv.org/abs/2503.22230>.

888 Guangming Sheng, Chi Zhang, Zilingfeng Ye, Xibin Wu, Wang Zhang, Ru Zhang, Yanghua Peng,
 889 Haibin Lin, and Chuan Wu. Hybridflow: A flexible and efficient rlhf framework. *arXiv preprint*
 890 *arXiv: 2409.19256*, 2024.

891 Mingyang Song, Mao Zheng, Zheng Li, Wenjie Yang, Xuan Luo, Yue Pan, and Feng Zhang. Fastcurl:
 892 Curriculum reinforcement learning with progressive context extension for efficient training r1-like
 893 reasoning models, 2025. URL <https://arxiv.org/abs/2503.17287>.

894 Sainbayar Sukhbaatar, Zeming Lin, Ilya Kostrikov, Gabriel Synnaeve, Arthur Szlam, and Rob
 895 Fergus. Intrinsic motivation and automatic curricula via asymmetric self-play, 2018. URL
 896 <https://arxiv.org/abs/1703.05407>.

897 Kimi Team, Angang Du, Bofei Gao, Bowei Xing, Changjiu Jiang, Cheng Chen, Cheng Li, Chenjun
 898 Xiao, Chenzhuang Du, Chonghua Liao, Chunling Tang, Congcong Wang, Dehao Zhang, Enming
 899 Yuan, Enzhe Lu, Fengxiang Tang, Flood Sung, Guangda Wei, Guokun Lai, Haiqing Guo, Han
 900 Zhu, Hao Ding, Hao Hu, Hao Yang, Hao Zhang, Haotian Yao, Haotian Zhao, Haoyu Lu, Haoze
 901 Li, Haozhen Yu, Hongcheng Gao, Huabin Zheng, Huan Yuan, Jia Chen, Jianhang Guo, Jianlin
 902 Su, Jianzhou Wang, Jie Zhao, Jin Zhang, Jingyuan Liu, Junjie Yan, Junyan Wu, Lidong Shi,
 903 Ling Ye, Longhui Yu, Mengnan Dong, Neo Zhang, Ningchen Ma, Qiwei Pan, Qucheng Gong,
 904 Shaowei Liu, Shengling Ma, Shupeng Wei, Sihan Cao, Siying Huang, Tao Jiang, Weihao Gao,
 905 Weimin Xiong, Weiran He, Weixiao Huang, Weixin Xu, Wenhao Wu, Wenyang He, Xianghui
 906 Wei, Xianqing Jia, Xingzhe Wu, Xinran Xu, Xinxing Zu, Xinyu Zhou, Xuehai Pan, Y. Charles,
 907 Yang Li, Yangyang Hu, Yangyang Liu, Yanru Chen, Yejie Wang, Yibo Liu, Yidao Qin, Yifeng Liu,
 908 Ying Yang, Yiping Bao, Yulun Du, Yuxin Wu, Yuzhi Wang, Zaida Zhou, Zhaoji Wang, Zhaowei
 909 Li, Zhen Zhu, Zheng Zhang, Zhexu Wang, Zhilin Yang, Zhiqi Huang, Zihao Huang, Ziyao Xu,
 910 Zonghan Yang, and Zongyu Lin. Kimi k1.5: Scaling reinforcement learning with llms, 2025. URL
 911 <https://arxiv.org/abs/2501.12599>.

912 RUCAIBox STILL Team. Still-3-1.5b-preview: Enhancing slow thinking abilities of small models
 913 through reinforcement learning. 2025. URL https://github.com/RUCAIBox/Slow_Thinking_with_LLMs.

918 Rui Wang, Joel Lehman, Jeff Clune, and Kenneth O. Stanley. Paired open-ended trailblazer (poet):
 919 Endlessly generating increasingly complex and diverse learning environments and their solutions,
 920 2019. URL <https://arxiv.org/abs/1901.01753>.

921 Yiping Wang, Qing Yang, Zhiyuan Zeng, Liliang Ren, Lucas Liu, Baolin Peng, Hao Cheng, Xuehai
 922 He, Kuan Wang, Jianfeng Gao, Weizhu Chen, Shuohang Wang, Simon Shaolei Du, and Yelong
 923 Shen. Reinforcement learning for reasoning in large language models with one training example,
 924 2025. URL <https://arxiv.org/abs/2504.20571>.

925 Liang Wen, Yunke Cai, Fenrui Xiao, Xin He, Qi An, Zhenyu Duan, Yimin Du, Junchen Liu, Lifu
 926 Tang, Xiaowei Lv, Haosheng Zou, Yongchao Deng, Shousheng Jia, and Xiangzheng Zhang.
 927 Light-r1: Curriculum sft, dpo and rl for long cot from scratch and beyond, 2025. URL <https://arxiv.org/abs/2503.10460>.

928 Yixin Ye, Zhen Huang, Yang Xiao, Ethan Chern, Shijie Xia, and Pengfei Liu. Limo: Less is more for
 929 reasoning, 2025. URL <https://arxiv.org/abs/2502.03387>.

930 Qiying Yu, Zheng Zhang, Ruofei Zhu, Yufeng Yuan, Xiaochen Zuo, Yu Yue, Tiantian Fan, Gaohong
 931 Liu, Lingjun Liu, Xin Liu, Haibin Lin, Zhiqi Lin, Bole Ma, Guangming Sheng, Yuxuan Tong, Chi
 932 Zhang, Mofan Zhang, Wang Zhang, Hang Zhu, Jinhua Zhu, Jiaze Chen, Jiangjie Chen, Chengyi
 933 Wang, Hongli Yu, Weinan Dai, Yuxuan Song, Xiangpeng Wei, Hao Zhou, Jingjing Liu, Wei-Ying
 934 Ma, Ya-Qin Zhang, Lin Yan, Mu Qiao, Yonghui Wu, and Mingxuan Wang. Dapo: An open-source
 935 llm reinforcement learning system at scale, 2025. URL <https://arxiv.org/abs/2503.14476>.

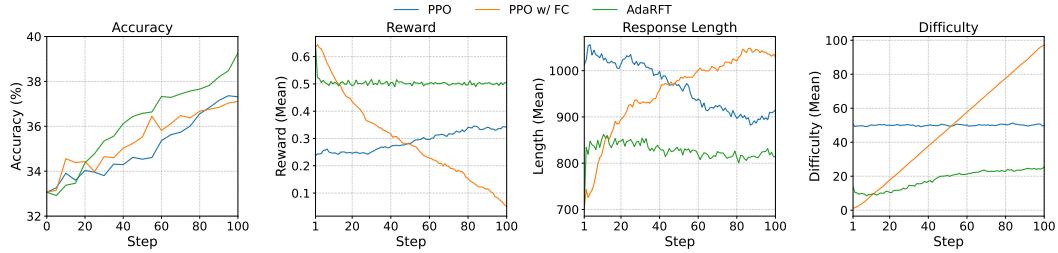
936 Wojciech Zaremba and Ilya Sutskever. Learning to execute, 2015. URL <https://arxiv.org/abs/1410.4615>.

937 Weihao Zeng, Yuzhen Huang, Qian Liu, Wei Liu, Keqing He, Zejun Ma, and Junxian He. Simplerl-
 938 zoo: Investigating and taming zero reinforcement learning for open base models in the wild, 2025.
 939 URL <https://arxiv.org/abs/2503.18892>.

940 Andrew Zhao, Yiran Wu, Yang Yue, Tong Wu, Quentin Xu, Yang Yue, Matthieu Lin, Shenzhi Wang,
 941 Qingyun Wu, Zilong Zheng, and Gao Huang. Absolute zero: Reinforced self-play reasoning with
 942 zero data, 2025. URL <https://arxiv.org/abs/2505.03335>.

943 Zyphra. Zr1-1.5b, a small but powerful reasoning model for math
 944 and code, 2025. URL <https://www.zyphra.com/post/introducing-zr1-1-5b-a-small-but-powerful-math-code-reasoning-model>.

945
 946
 947
 948
 949
 950
 951
 952
 953
 954
 955
 956
 957
 958
 959
 960
 961
 962
 963
 964
 965
 966
 967
 968
 969
 970
 971

972 A FURTHER DISCUSSION
973
974975 A.1 TRAINING DYNAMICS: ADARFT VS. FIXED CURRICULUM VS. PPO
976
977

988 Figure 5: Training dynamics of standard PPO, PPO with a fixed curriculum (PPO w/ FC), and
989 ADARFT on the Qwen 2.5 Math 1.5B model under the uniform data distribution. We plot accuracy,
990 reward, response length, and average difficulty of sampled training problems over training steps.
991 Curves are exponentially smoothed ($\alpha = 0.3$) for clarity.

992
993
994 To further contextualize the role of adaptive curricula, we analyze the training dynamics of fixed
995 curriculum (PPO w/ FC), standard PPO, and ADARFT when training the Qwen 2.5 Math 1.5B model
996 on the uniform distribution. Figure 5 illustrates the evolution of accuracy, reward, response length,
997 and sampled difficulty across the first 100 training steps.

998 The accuracy curves highlight the central tradeoff of fixed curricula. PPO w/ FC initially converges
999 slightly faster than standard PPO, benefiting from early exposure to easier problems. However, it
1000 ultimately underperforms compared to ADARFT because its difficulty schedule increases indepen-
1001 dently of the model’s actual learning progress. As a result, the model is pushed into harder problem
1002 regimes too quickly. This misalignment is clearly visible in the reward curve. Early in training, PPO
1003 w/ FC achieves rewards well above 0.5, indicating that the model is initially exposed to problems that
1004 are easier than its current capability. However, as training progresses, the fixed curriculum increases
1005 difficulty at a rate that outpaces the model’s learning speed. As a result, the reward drops below
1006 0.5 and continues declining, showing that the model is increasingly confronted with problems it
1007 cannot yet solve. This mismatch limits the effectiveness of the updates and ultimately leads to slower
1008 convergence compared to an adaptive curriculum.

1009 The response length patterns reinforce this interpretation. Because longer responses typically corre-
1010 spond to more difficult reasoning tasks, the rapid increase in response length under PPO w/ FC shows
1011 that the curriculum escalates difficulty faster than the model can adapt. In contrast, standard PPO
1012 maintains more stable lengths but lacks the structured progression necessary for efficient learning.
1013 ADARFT, by comparison, keeps response lengths moderate and gradually increasing, consistent with
1014 its difficulty traces: the curriculum raises problem difficulty only when the model’s reward stays near
1015 the target value, ensuring that the model always receives examples that are challenging but solvable.

1016 This particular dynamic suggest a general underlying issue: since the model’s perceived difficulty
1017 changes over the course of training, a fixed curriculum cannot remain aligned, and different model-
1018 dataset combinations may experience extended periods of being over-challenged or under-challenged.
1019 This mismatch limits the effectiveness of updates and leads to slower convergence compared to
1020 an adaptive curriculum. In contrast, ADARFT is designed to handle precisely this challenge: by
1021 adjusting the difficulty schedule based on the model’s reward signal, it continually matches training
1022 difficulty to the model’s evolving capability, ensuring sustained learning progress throughout training.
1023 Without the ability to adjust to real-time model performance, fixed schedules risk either overwhelming
1024 the model or wasting compute on overly easy tasks. ADARFT avoids both extremes by maintaining
1025 the average reward near 0.5, automatically pacing the introduction of harder problems as the model
1026 improves. This adaptive alignment between task difficulty and model capability leads to smoother
1027 reward trajectories, controlled response lengths, and ultimately more stable and efficient learning.

1026
1027

A.2 DATA DIFFICULTY ON MODEL PERFORMANCE

1028
1029
1030
1031
1032
1033
1034
1035
1036
1037
1038
1039
1040

To better understand the effect of data difficulty on model performance, we introduce two additional data distributions: easy-extreme and hard-extreme. Unlike the skew-difficult and skew-easy distributions, which still include a mix of difficulty levels, the easy-extreme and hard-extreme sets consist exclusively of the most polarized examples. Specifically, easy-extreme contains only the easiest samples with difficulty levels no greater than 15, while hard-extreme includes only the hardest samples with difficulty levels of at least 97. Each of these extreme distributions consists of approximately 8,000 samples, providing a focused and controlled evaluation of model behavior under minimal or maximal difficulty conditions. We trained a Qwen 2.5 7B model on each of the two extreme distributions using PPO, and compared their performance to models trained on the uniform distribution with PPO (Uniform) and with ADARFT instantiated with PPO (Uniform + ADARFT), as described in Section 5. The results are presented in Figure 6. The key takeaway is that training on only overly easy or hard problems fails to provide useful learning signals, reinforcing the need for ADARFT to adaptively steer models toward challenges matched to their current ability.

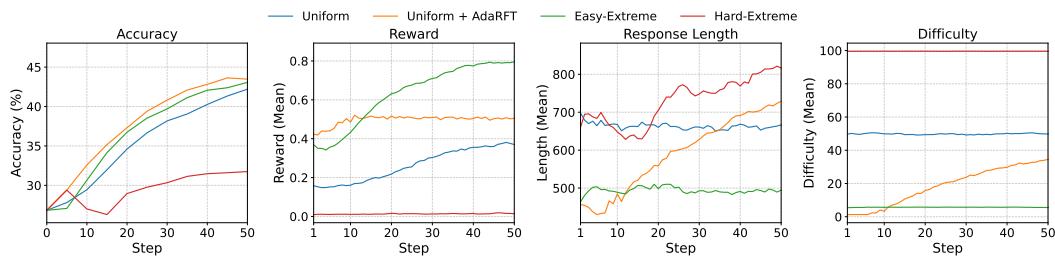
1041
1042
1043
1044
1045
1046
1047
1048
10491050
1051
1052

Figure 6: Performance comparison of Qwen 2.5 7B trained on different data distributions using PPO (Uniform, Easy-Extreme, Hard-Extreme) and ADARFT instantiated with PPO (Uniform + ADARFT). For clarity, curves are exponentially smoothed ($\alpha = 0.3$) to reduce noise.

1053
1054
1055
1056
1057
1058
1059
1060
1061

Accuracy. The leftmost panel of Figure 6 shows that uniform + ADARFT achieves the highest overall accuracy throughout training, outperforming both uniform and the two extreme settings. This highlights the effectiveness of ADARFT in guiding the model through an optimal difficulty progression. In contrast, hard-extreme struggles significantly, with a flat and lower trajectory, indicating that exposing the model only to very difficult problems limits learning progress. This suggests that without a gradual exposure strategy, models trained on only the hardest problems are unable to bootstrap their capabilities effectively.

1062
1063
1064
1065
1066
1067
1068
1069
1070
1071
1072
1073
1074
1075
1076
1077

Reward. The reward trends provide important clues about learning dynamics. The easy-extreme setup achieves the fastest reward improvement during early training, surpassing both uniform and hard-extreme. In particular, easy-extreme consistently operates in a reward range between 0.4 and 0.6 during early training, which corresponds to a success rate that is both challenging and attainable. In contrast, the reward of the uniform and hard-extreme setup lingers below 0.2 in early training, leading to slower learning. This suggests that training on problems with intermediate difficulty—those that are neither trivially easy nor prohibitively hard—provides the most effective learning signal. Notably, ADARFT is explicitly designed to exploit this insight: by setting the target reward $\beta = 0.5$, we encourage the model to train on problems that match this “productive struggle” zone. As shown by the uniform + ADARFT curve, the algorithm successfully maintains an average reward near 0.5 throughout training, allowing the model to learn at an optimal pace. Notably, while the uniform setup eventually reaches a reward of nearly 0.5 by step 50, it does not result in faster learning. This is likely because the model is already fairly well trained by that stage, so the additional reward signal contributes less to further improvement. In contrast, the hard-extreme model receives almost no reward signal for most of the training, while the uniform setup shows slower and more gradual reward accumulation.

1078
1079

Response Length. The response length panel reveals how the complexity of generated solutions evolves during training. The hard-extreme model consistently produces the longest responses, with length increasing steadily, reflecting the higher complexity and reasoning depth required by the

1080 hardest problems. In contrast, the easy-extreme setup maintains short and stable responses, consistent
 1081 with its simpler problem set. The uniform and uniform + ADARFT setups fall between these two
 1082 extremes. Notably, uniform + ADARFT shows a gradual increase in response length over time. This
 1083 trend aligns with the behavior of the curriculum learning algorithm: as the model improves, it is
 1084 exposed to increasingly difficult problems, which naturally demand more elaborate reasoning and
 1085 longer solutions. This dynamic suggests that response length can serve as a useful proxy for problem
 1086 difficulty and reasoning complexity during training.

1087
 1088 **Difficulty.** Finally, the difficulty panel illustrates how problem difficulty evolves under each setup.
 1089 The easy-extreme and hard-extreme curves remain flat, confirming that these datasets contain only
 1090 problems from the tail ends of the difficulty spectrum (i.e., ≤ 15 and ≥ 97 , respectively). The
 1091 uniform curve is centered around 50, as expected, while uniform + ADARFT shows a steady increase
 1092 in difficulty over time. This adaptive progression confirms that curriculum learning effectively steers
 1093 the model from easier to harder problems, aligning difficulty with the model’s evolving capabilities.

1094 A.3 ADARFT WITH DIVERSE RL ALGORITHMS

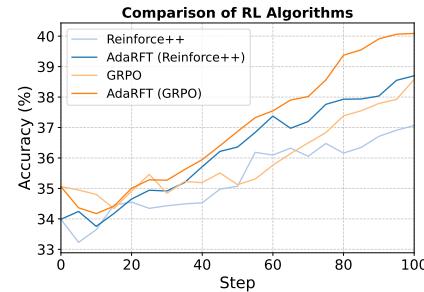
1095 To evaluate the generality of ADARFT beyond PPO, we
 1096 trained the Qwen 2.5 Math 1.5B model on a skew-difficult
 1097 data distribution using two alternative reinforcement learn-
 1098 ing algorithms: REINFORCE++ and GRPO (see imple-
 1099 mentation details in Appendix B). As shown in Figure 7,
 1100 ADARFT significantly improves both the convergence
 1101 speed and final accuracy across these variants. Across
 1102 both cases, the adaptive curriculum acts orthogonally to
 1103 the underlying optimization method. These results rein-
 1104 force the plug-and-play nature of ADARFT: it consistently
 1105 enhances sample efficiency and policy robustness across al-
 1106 gorithmic choices, making it broadly applicable in diverse
 1107 reinforcement finetuning pipelines. Notably, this general-
 1108 ization holds without any additional tuning or algorithm-
 1109 specific modifications, underscoring the practical utility of
 1110 curriculum-aware training in both lightweight and computationally-heavy RFT settings.

1112 A.4 TRAINING ON LLM-ESTIMATED DIFFICULTY

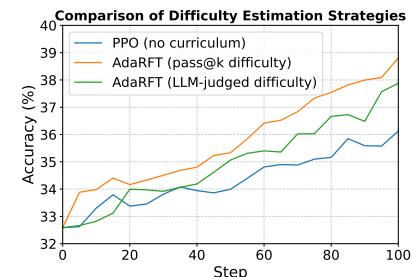
1113 In addition to rollout-based difficulty estimation, we explore an alternative strategy that uses LLM-
 1114 judged difficulty levels to guide curriculum construction. As described in Section 4.1, we prompt
 1115 GPT-4o (gpt-4o-0806) to assign difficulty levels to math problems in the DeepScaleR dataset
 1116 according to the AoPS rubric. This approach offers a lightweight and scalable alternative to computing
 1117 pass@ k success rates from model rollouts, making it especially attractive in low-resource scenarios.

1118 To assess the effectiveness of this strategy, we train a Qwen
 1119 2.5 Math 1.5B model on the skew-difficult distribution using
 1120 ADARFT (PPO) with two curriculum schedules: one
 1121 based on rollout-derived pass@ k difficulty, and the other
 1122 guided by GPT-4o’s difficulty ratings. Since the LLM-
 1123 judged difficulty is on a scale of 1 to 5 (rather than 0 to
 1124 100), we set the step size hyperparameter $\eta = 2.5$ to align
 1125 the difficulty adjustment magnitude with the reward sig-
 1126 nals. All other hyperparameters are kept unchanged. As
 1127 shown in Figure 8, both curriculum strategies outperform
 1128 standard PPO without curriculum learning. While rollout-
 1129 based difficulty estimation yields the strongest gains, the
 1130 LLM-judged curriculum still provides a noticeable im-
 1131 provement over the baseline.

1132 These results demonstrate that ADARFT remains effective even when the difficulty signal is derived
 1133 from heuristic or approximate sources like LLM judgments. Although less precise than empirical



1134 Figure 7: Comparison between models
 1135 trained with and without AdaRFT using
 1136 REINFORCE++ and GRPO.



1137 Figure 8: Comparison of different diffi-
 1138 culty estimation strategies.

1134 pass@ k metrics, the LLM-based difficulty still provides enough structure to enable meaningful
 1135 curriculum adaptation. This makes it a practical fallback when rollout computation is too costly, and
 1136 suggests that future work could explore hybrid approaches that combine lightweight heuristics with
 1137 periodic empirical calibration.

1139 B IMPLEMENTATION DETAILS

1141 B.1 TRAINING CONFIGURATION

1143 We trained both the actor and critic models using the PPO algorithm on a single node with 8 A100
 1144 GPUs. Each model was trained for approximately 100 optimization steps using the `veRL` library
 1145 (Sheng et al., 2024). We used two model variants: `Qwen2.5-7B` and `Qwen2.5-MATH-1.5B`.
 1146 The latter has a shorter context window, so we adjusted the max response length and the sequence
 1147 parallel size accordingly.

1148 Table 3 summarizes the core hyperparameter settings used across all three algorithms: PPO, GRPO,
 1149 and REINFORCE++. We highlight both shared defaults and algorithm-specific overrides, including
 1150 KL treatment modes, rollout settings, and critic configurations.

1152 B.2 DERIVING THE TARGET-DIFFICULTY UPDATE RULE FROM A LINEAR MAPPING

1154 A central component of our curriculum mechanism is the update of the target difficulty T based on
 1155 the model’s observed reward performance. While the final update rule (Eq. 13) involves a hyperbolic
 1156 tangent, it is in fact a smooth and stabilized variant of a standard linear mapping between reward
 1157 space and difficulty space. We derive it here for clarity.

1158 **Step 1: Linear mapping between two intervals.** The classical linear interpolation formula for
 1159 mapping a value $v \in [x, y]$ to a target interval $[a, b]$ is

$$1161 v' = a + \frac{(v - x)(b - a)}{y - x}. \quad (5)$$

1163 If we directly map the average reward $R_{\text{avg}} \in [r_{\min}, r_{\max}]$ to the difficulty range $[d_{\min}, d_{\max}]$, we
 1164 obtain

$$1165 T_{\text{naive}}(R_{\text{avg}}) = d_{\min} + \frac{(R_{\text{avg}} - r_{\min})(d_{\max} - d_{\min})}{r_{\max} - r_{\min}}. \quad (6)$$

1167 In our main setting, $r_{\min} = 0$, $r_{\max} = 1$, and $[d_{\min}, d_{\max}] = [0, 100]$, so the naive mapping simplifies
 1168 to

$$1169 T_{\text{naive}}(R_{\text{avg}}) = 100 R_{\text{avg}}. \quad (7)$$

1171 **Step 2: Mapping reward deviation instead of absolute reward.** For curriculum learning, we do
 1172 not wish to reassign a new difficulty level at every step. Instead, we aim to *adjust* the current target
 1173 difficulty depending on whether the model is performing above or below a desired target success rate
 1174 β . We therefore consider the deviation

$$1175 \delta = R_{\text{avg}} - \beta. \quad (8)$$

1177 Given $R_{\text{avg}} \in [0, 1]$, the deviation satisfies $\delta \in [r_{\min} - \beta, r_{\max} - \beta] = [-\beta, 1 - \beta]$. With the
 1178 common choice $\beta = 0.5$, this becomes $\delta \in [-0.5, 0.5]$.

1179 Applying the linear mapping rule equation 5 from the deviation range $[-0.5, 0.5]$ to a symmetric
 1180 difficulty-change interval $[-\Delta, \Delta]$ yields

$$1181 \Delta T_{\text{lin}}(\delta) = -\Delta + \frac{(\delta - (-0.5))(\Delta - (-\Delta))}{0.5 - (-0.5)} = 2\Delta \delta. \quad (9)$$

1184 Thus the naive linear controller becomes

$$1185 T'_{\text{lin}} = T + 2\Delta (R_{\text{avg}} - \beta). \quad (10)$$

1187 This already captures the desired behavior: difficulty increases when performance exceeds the target,
 decreases when performance falls short, and remains stable when $R_{\text{avg}} = \beta$.

1188
 1189 **Step 3: Stabilizing the update via a smooth saturating nonlinearity.** A purely linear controller
 1190 may cause excessively large changes when the reward deviation is large or noisy. To obtain a stable
 1191 update rule, we replace the linear term with a smooth, odd, saturating nonlinearity. The hyperbolic
 1192 tangent is a natural choice: it behaves linearly near zero (which recovers the linear mapping) and
 1193 saturates as its argument grows.

1194 We therefore define a smoothed difficulty adjustment

$$1195 \Delta T(\delta) = \eta \cdot \tanh(\alpha(R_{\text{avg}} - \beta)). \quad (11)$$

1196
 1197 Here, η sets the maximum update magnitude and α controls the sensitivity around the target reward.
 1198 For small deviations, $\tanh(z) \approx z$, so locally

$$1199 \Delta T(\delta) \approx \eta\alpha(R_{\text{avg}} - \beta), \quad (12)$$

1200
 1201 recovering a linear controller with effective slope $\eta\alpha$ while ensuring global boundedness.

1202
 1203 **Step 4: Clipping to the valid difficulty range.** To ensure the target difficulty remains within the
 1204 observed range of the data, we apply a final clipping:

$$1205 \Delta T' = \text{clip}(\Delta T + \eta \cdot \tanh(\alpha(R_{\text{avg}} - \beta)), d_{\min}, d_{\max}). \quad (13)$$

1206
 1207 The full update rule equation 13 is therefore a direct, smoothed generalization of a naive linear
 1208 mapping between reward deviations and difficulty adjustments. It preserves the intuitive behavior
 1209 of the linear controller near the target reward, while the saturating nonlinearity and clipping ensure
 1210 stable, bounded, and data-consistent curriculum updates.

1211
 1212 Because the reward is bounded in $[0, 1]$ and the difficulty metric spans $[0, 100]$, we set the step size
 1213 $\eta = 50$ to align their scales. The modulation parameter $\alpha = 2$ ensures smooth and controlled
 1214 progression throughout training.

1215 B.3 PROMPT FOR DIFFICULTY ESTIMATION USING LLM AS A JUDGE

1216
 1217 The prompt used for difficulty estimation (as described in Section 4.1) is shown in Table 4, Table 5,
 1218 and Table 6. The descriptions of the difficulty scales and examples are sourced from the AoPS Wiki.²
 1219 Although GPT-4o is prompted to rate problem difficulty on a scale from 1 to 10, we found that over
 1220 95% of the problems fall within the range of 1 to 5. Therefore, we clip the scores and use a revised
 1221 scale from 1 to 5. In addition to integer scores, we also allow half-point increments such as 1.5, 2.0,
 1222 and 2.5 for finer-grained difficulty estimation.

1223 C THE USE OF LARGE LANGUAGE MODELS FOR ICLR 2026

1224
 1225 In this ICLR submission, large language models (LLMs) were used solely as writing aids for grammar
 1226 correction, wording refinement, and text polishing. They were not employed for idea generation,
 1227 technical contributions, or any aspect of the research beyond enhancing readability and clarity.

1228
 1229
 1230
 1231
 1232
 1233
 1234
 1235
 1236
 1237
 1238
 1239
 1240

1241 ²https://artofproblemsolving.com/wiki/index.php/AoPS_Wiki:Competition_ratings

1242
 1243
 1244
 1245
 1246
 1247
 1248
 1249
 1250

Table 3: Comparison of training hyperparameters for PPO, GRPO, and REINFORCE++ using the veRL library. Shared defaults are used unless overridden.

| Category | Parameter | PPO | GRPO | REINFORCE++ |
|--|---------------------------------|------------------------|--------------------|--------------------|
| <i>Algorithm-Specific Settings</i> | | | | |
| General | Advantage estimator | GAE | GRPO | REINFORCE++ |
| | Gamma (γ) | 1.0 | — | — |
| | Lambda (λ) | 1.0 | — | — |
| | Batch size | 1024 | 1024 | 1024 |
| | Max prompt length | 1024 | 1024 | 1024 |
| | Gradient checkpointing | Enabled | Enabled | Enabled |
| Actor | Learning rate | 1×10^{-6} | 1×10^{-6} | 1×10^{-6} |
| | Mini-batch size | 1024 | 1024 | 1024 |
| | Dynamic batch size | Enabled | Enabled | Enabled |
| | KL penalty role | Reward | Loss | Loss |
| | KL loss type | Fixed | Low-variance KL | MSE |
| | KL loss coefficient (β) | 0.001 | 0.001 | 0.001 |
| | Entropy coefficient | 0.001 | 0.001 | 0 |
| | Clip ratio | 0.2 | 0.2 | 0.2 |
| | Gradient clipping | 1.0 | 1.0 | 1.0 |
| | Sequence parallel size | Model-specific | Model-specific | Model-specific |
| Rollout | Backend | vLLM | vLLM | vLLM |
| | Tensor model parallel size | 2 | 2 | 2 |
| | Rollouts per sample | 1 | 8 | 1 |
| | Nucleus sampling p | 1.0 | 1.0 | 1.0 |
| | GPU memory utilization | 0.5 | 0.5 | 0.5 |
| | Sampling temperature | 1.0 | 1.0 | 1.0 |
| Critic | Warmup steps | 0 | — | — |
| | Learning rate | 1×10^{-5} | — | — |
| | Sequence parallel size | Model-specific | — | — |
| <i>Model-Specific Overrides (shared across all algorithms)</i> | | | | |
| Qwen2.5-7B | Max response length | 8000 | 8000 | 8000 |
| | Sequence parallel size | 2 | 2 | 2 |
| | Max token length / GPU | 8000 | 8000 | 8000 |
| | Max response length | 3000 | 3000 | 3000 |
| | Qwen2.5-MATH-1.5B | Sequence parallel size | 1 | 1 |
| Curriculum Learning | Max token length / GPU | 16000 | 16000 | 16000 |
| | <i>ADARFT Parameters</i> | | | |
| | Target reward (β) | 0.5 | 0.5 | 0.5 |
| | Sensitivity (α) | 2 | 2 | 2 |
| | Step size (η) | 50 | 50 | 50 |
| | Initial difficulty (T) | 0 | 0 | 0 |

1288
 1289
 1290
 1291
 1292
 1293
 1294
 1295

1296

1297

1298

1299

1300

1301 **Prompt for Difficulty Estimation (Part 1)**

1302

Math Problem

{problem}

1305

Your Task

You are a subject matter expert in mathematics tasked with evaluating the difficulty level of individual math problems. Your assessment should be objective and based on a detailed difficulty scale provided below. Your judgment will help calibrate and categorize problems for use in educational settings or assessments. Be thorough, fair, and consistent in your evaluation.

1311

Difficulty Scale

1: Problems strictly for beginner, on the easiest elementary school or middle school levels (MOEMS, MATHCOUNTS School, AMC 8 1-10, AMC 10 1-10, easier AMC 12 1-5, and others that involve standard techniques introduced up to the middle school level), most traditional middle/high school word problems.

1.5: Problems for stronger beginner students, on the level of the middling problems in most middle school contests (AMC 8 11-20, harder AMC 10 1-10, AMC 12 1-5, and those others that force students to apply their school-level knowledge to slightly more challenging problems), traditional middle/high school word problems with more complex problem solving.

2: For motivated beginners, harder questions from the previous categories (AMC 8 21-25, MATHCOUNTS Chapter (Sprint 21-30, Target 6-8), MATHCOUNTS States/Nationals, AMC 10 11-15, AMC 12 5-10, easiest AIME 1-3)

2.5: More advanced beginner problems, hardest questions from previous categories (Harder AMC 8 21-25, harder MATHCOUNTS States questions, AMC 10 16-20, AMC 12 11-15, usual AIME 1-3)

3: Early intermediate problems that require more creative thinking (harder MATHCOUNTS National questions, AMC 10 21-25, AMC 12 15-20, hardest AIME 1-3, usual AIME 4-6).

4: Intermediate-level problems (AMC 12 21-25, hardest AIME 4-6, usual AIME 7-10).

5: More difficult AIME problems (11-13), simple proof-based Olympiad-style problems (early JBMO questions, easiest USAJMO 1/4).

6: High-leveled AIME-style questions (14/15). Introductory-leveled Olympiad-level questions (harder USAJMO 1/4 and easier USAJMO 2/5, easier USAMO and IMO 1/4).

7: Tougher Olympiad-level questions, may require more technical knowledge (harder USAJMO 2/5 and most USAJMO 3/6, extremely hard USAMO and IMO 1/4, easy-medium USAMO and IMO 2/5).

8: High-level Olympiad-level questions (medium-hard USAMO and IMO 2/5, easiest USAMO and IMO 3/6).

9: Expert Olympiad-level questions (average USAMO and IMO 3/6).

9.5: The hardest problems appearing on Olympiads which the strongest students could reasonably solve (hard USAMO and IMO 3/6).

10: Historically hard problems, generally unsuitable for very hard competitions (such as the IMO) due to being exceedingly tedious, long, and difficult (e.g. very few students are capable of solving on a worldwide basis).

1343

1344

1345

1346

1347

1348

1349

Table 4: Prompt for difficulty estimation using LLM as a judge.

1350

1351 **Prompt for Difficulty Estimation (Part 2)**

1352

1353 **# Examples**

1354 For reference, here are some sample problems from each of the difficulty levels 1-10:

1355 <1: Jamie counted the number of edges of a cube, Jimmy counted the numbers of corners, and Judy counted the number of faces. They then added the three numbers. What was the resulting sum? (2003 AMC 8, Problem 1)

1356 1: How many integer values of x satisfy $|x| < 3\pi$? (2021 Spring AMC 10B, Problem 1)

1357 1.5: A number is called flippy if its digits alternate between two distinct digits. For example, 2020 and 37373 are flippy, but 3883 and 123123 are not. How many five-digit flippy numbers are divisible by 15? (2020 AMC 8, Problem 19)

1358 2: A fair 6-sided die is repeatedly rolled until an odd number appears. What is the probability that every even number appears at least once before the first occurrence of an odd number? (2021 Spring AMC 10B, Problem 18)

1359 2.5: A , B , C are three piles of rocks. The mean weight of the rocks in A is 40 pounds, the mean weight of the rocks in B is 50 pounds, the mean weight of the rocks in the combined piles A and B is 43 pounds, and the mean weight of the rocks in the combined piles A and C is 44 pounds. What is the greatest possible integer value for the mean in pounds of the rocks in the combined piles B and C ? (2013 AMC 12A, Problem 16)1360 3: Triangle ABC with $AB = 50$ and $AC = 10$ has area 120. Let D be the midpoint of \overline{AB} , and let E be the midpoint of \overline{AC} . The angle bisector of $\angle BAC$ intersects \overline{DE} and \overline{BC} at F and G , respectively. What is the area of quadrilateral $FDBG$? (2018 AMC 10A, Problem 24)1361 3.5: Find the number of integer values of k in the closed interval $[-500, 500]$ for which the equation $\log(kx) = 2\log(x+2)$ has exactly one real solution. (2017 AIME II, Problem 7)1362 4: Define a sequence recursively by $x_0 = 5$ and

1363
$$x_{n+1} = \frac{x_n^2 + 5x_n + 4}{x_n + 6}$$

1364 for all nonnegative integers n . Let m be the least positive integer such that

1365
$$x_m \leq 4 + \frac{1}{2^{20}}.$$

1366 In which of the following intervals does m lie?1367 (A) $[9, 26]$ (B) $[27, 80]$ (C) $[81, 242]$ (D) $[243, 728]$ (E) $[729, \infty)$ (2019

1368 AMC 10B, Problem 24 and 2019 AMC 12B, Problem 22)

1369 4.5: Find, with proof, all positive integers n for which $2^n + 12^n + 2011^n$ is a perfect square. (USAJMO 2011/1)1370 5: Find all triples (a, b, c) of real numbers such that the following system holds:

1371
$$a + b + c = \frac{1}{a} + \frac{1}{b} + \frac{1}{c},$$

1372
$$a^2 + b^2 + c^2 = \frac{1}{a^2} + \frac{1}{b^2} + \frac{1}{c^2}.$$

1373 (JBMO 2020/1)

1374 5.5: Triangle ABC has $\angle BAC = 60^\circ$, $\angle CBA \leq 90^\circ$, $BC = 1$, and $AC \geq AB$. Let H , I , and O be the orthocenter, incenter, and circumcenter of $\triangle ABC$, respectively. Assume that the area of pentagon $BCOIH$ is the maximum possible. What is $\angle CBA$? (2011 AMC 12A, Problem 25)1375 6: Let $\triangle ABC$ be an acute triangle with circumcircle ω , and let H be the intersection of the altitudes of $\triangle ABC$. Suppose the tangent to the circumcircle of $\triangle HBC$ at H intersects ω at points X and Y with $HA = 3$, $HX = 2$, and $HY = 6$. The area of $\triangle ABC$ can be written in the form $m\sqrt{n}$, where m and n are positive integers, and n is not divisible by the square of any prime. Find $m + n$. (2020 AIME I, Problem 15)

1402

1403

Table 5: Prompt for difficulty estimation using LLM as a judge.

1404

1405

1406

1407

1408

1409

1410

1411

1412

1413

1414

1415

1416

1417

1418

1419

1420

1421

1422

1423

1424

1425

1426

1427

1428

1429

1430

1431

1432

1433

1434

1435

1436

1437

1438

1439

1440

1441

1442

1443

1444

1445

1446

1447

1448

1449

1450

1451

1452

1453

1454

1455

1456

1457

Prompt for Difficulty Estimation (Part 3)

6.5: Rectangles BCC_1B_2 , CAA_1C_2 , and ABB_1A_2 are erected outside an acute triangle ABC . Suppose that

$$\angle BC_1C + \angle CA_1A + \angle AB_1B = 180^\circ.$$

Prove that lines B_1C_2 , C_1A_2 , and A_1B_2 are concurrent. (USAMO 2021/1, USAJMO 2021/2)
 7: We say that a finite set \mathcal{S} in the plane is balanced if, for any two different points A, B in \mathcal{S} , there is a point C in \mathcal{S} such that $AC = BC$. We say that \mathcal{S} is centre-free if for any three points A, B, C in \mathcal{S} , there is no point P in \mathcal{S} such that $PA = PB = PC$.

Show that for all integers $n \geq 3$, there exists a balanced set consisting of n points. Determine all integers $n \geq 3$ for which there exists a balanced centre-free set consisting of n points. (IMO 2015/1)

7.5: Let \mathbb{Z} be the set of integers. Find all functions $f : \mathbb{Z} \rightarrow \mathbb{Z}$ such that

$$xf(2f(y) - x) + y^2f(2x - f(y)) = \frac{f(x)^2}{x} + f(yf(y))$$

for all $x, y \in \mathbb{Z}$ with $x \neq 0$. (USAMO 2014/2)

8: For each positive integer n , the Bank of Cape Town issues coins of denomination $\frac{1}{n}$. Given a finite collection of such coins (of not necessarily different denominations) with total value at most $99 + \frac{1}{2}$, prove that it is possible to split this collection into 100 or fewer groups, such that each group has total value at most 1. (IMO 2014/5)

8.5: Let I be the incentre of acute triangle ABC with $AB \neq AC$. The incircle ω of ABC is tangent to sides BC, CA , and AB at D, E , and F , respectively. The line through D perpendicular to EF meets ω at R . Line AR meets ω again at P . The circumcircles of triangle PCE and PBF meet again at Q .

Prove that lines DI and PQ meet on the line through A perpendicular to AI . (IMO 2019/6)

9: Let k be a positive integer and let S be a finite set of odd prime numbers. Prove that there is at most one way (up to rotation and reflection) to place the elements of S around the circle such that the product of any two neighbors is of the form $x^2 + x + k$ for some positive integer x . (IMO 2022/3)

9.5: An anti-Pascal triangle is an equilateral triangular array of numbers such that, except for the numbers in the bottom row, each number is the absolute value of the difference of the two numbers immediately below it. For example, the following is an anti-Pascal triangle with four rows which contains every integer from 1 to 10.

$$\begin{array}{cccc} & & & 4 \\ & & 2 & 6 \\ & 5 & 7 & 1 \\ 8 & 3 & 10 & 9 \end{array}$$

Does there exist an anti-Pascal triangle with 2018 rows which contains every integer from 1 to $1 + 2 + 3 + \dots + 2018$? (IMO 2018/3)

10: Prove that there exists a positive constant c such that the following statement is true: Consider an integer $n > 1$, and a set \mathcal{S} of n points in the plane such that the distance between any two different points in \mathcal{S} is at least 1. It follows that there is a line ℓ separating \mathcal{S} such that the distance from any point of \mathcal{S} to ℓ is at least $cn^{-1/3}$.

(A line ℓ separates a set of points S if some segment joining two points in S crosses ℓ .) (IMO 2020/6)

Return format

Please return the corresponding difficulty scale (integer) in $\backslash \text{box}\{ \}$

Table 6: Prompt for difficulty estimation using LLM as a judge.