

000 SQUEEZE THE SOAKED SPONGE: EFFICIENT OFF- 001 POLICY RFT FOR LARGE LANGUAGE MODEL 002 003 004

005 **Anonymous authors**

006 Paper under double-blind review

007 008 ABSTRACT 009

010 Reinforcement Learning (RL) has demonstrated its potential to improve the reasoning ability of Large Language Models (LLMs), yet most existing Reinforcement Finetuning (RFT) methods are inherently *on-policy* RL, failing to reuse historical data and thus preventing efficient scaling. In this work, we explore the potential of *off-policy* RL to leverage historical data for rollout-efficient RFT. Specifically, we propose **Reincarnating Mix**-policy Proximal Policy Gradient (**ReMix**), which enables on-policy RFT methods to leverage off-policy data. ReMix consists of three major components: (1) Mix-policy proximal policy gradient with an increased Update-To-Data (UTD) ratio that utilizes the data from both current and past policies for efficient training; (2) KL-Convex policy constraint that combines the KL constraints on the base and precedent model to balance stability and flexibility; (3) Policy reincarnation that replaces the base model with the mix-policy RFT model in the mid way of training and restarts on-policy training, to achieve a seamless transition from early efficiency to steady convergence. In our experiments, we train a series of ReMix models based on PPO, GRPO from 1.5B, 7B base models. On five math reasoning benchmarks (i.e., AIME’24, AMC’23, Minerva, OlympiadBench, and MATH500), ReMix achieves an average Pass@1 accuracy of **52.10%** (with **0.079M rollouts**) and **64.39%** (with **0.011M rollouts**) on 1.5B and 7B models, respectively. Compared with 15 recent advanced models, ReMix shows SOTA-level performance with an over **30x to 450x reduction in training cost in terms of rollout data volume**, demonstrating superior training efficiency. Additionally, our multifaceted analysis reveals insightful findings, including the implicit preference for shorter responses of off-policy RFT, the collapse mode of self-reflection under severe off-policyness, etc.

034 035 1 INTRODUCTION

036 The emergence of Large Language Models (LLMs) has lifted artificial intelligence to a next level, with the milestone works like (OpenAI, 2022; Jaech et al., 2024; Bai et al., 2022a; Trung et al., 2024; Guo et al., 2025). Consistent efforts are being made to push forward the limits of LLMs in performing deeper thinking and solving more complex tasks (Li et al., 2025b). Recently, Large Reasoning Models (LRMs) (Jaech et al., 2024; Guo et al., 2025; Kimi et al., 2025; Yang et al., 2025a) have taken the stage and attracted great attention, showing that a significant improvement of problem-solving ability can be achieved by a long human-like reasoning process (i.e., *slow thinking*), especially in scenarios like Math, Coding, Scientific Q&A, etc. One of the central recipes of LRM is Reinforcement Finetuning (RFT) (Trung et al., 2024). By treating the LLM as a policy model, the LLM can follow the philosophy of Reinforcement Learning (RL) (Sutton & Barto, 1998) and learn to reason and answer the queries according to the reward signals, e.g., either from a verifiable reward function (Guo et al., 2025) or a learned reward model (Bai et al., 2022a).

048 Although RFT opens another space for more powerful reasoning ability of LLMs, the longstanding and notorious shortcoming of RL — *sample inefficiency* — still exists. In another word, RFT 049 usually needs significantly more computational cost (e.g., rollouts and training) than SFT due to 050 its *trial-and-error* nature. The inefficiency of RL poses a stringent bottleneck on time and cost, 051 consequently preventing further scaling of model size and response length of LLMs. Currently, policy 052 gradient algorithms like PPO (Schulman et al., 2017), GRPO (Shao et al., 2024), RLOO (Ahmadian 053 et al., 2024) are widely adopted for RFT of LLMs due to their stable learning performance and

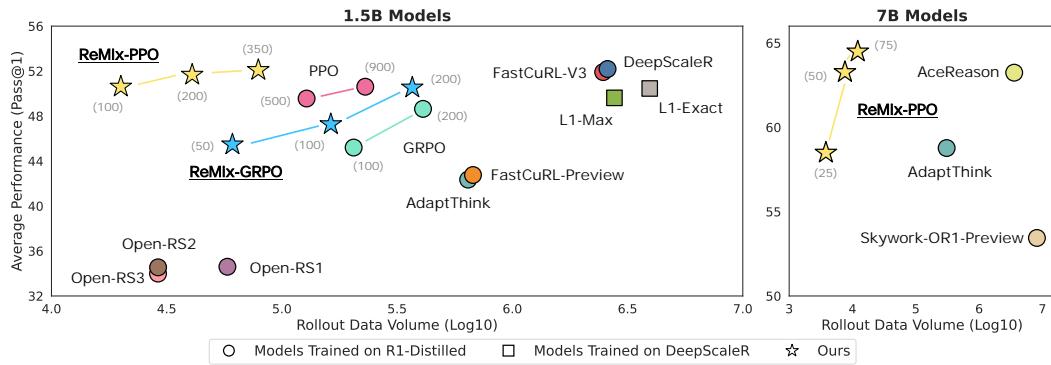


Figure 1: **Efficiency-performance comparison** for 1.5B models (*left*) and 7B models (*right*) in terms of Rollout Data Volume (i.e., total number of responses generated during training) v.s., Average Pass@1 Accuracy on five math reasoning benchmarks. An ideal model should appear in the top-left.

friendliness to engineering. However, they are all *on-policy* algorithms, which are known to be sample inefficient as the data generated by the online policy is dropped after each iteration. In the literature of RL, *off-policy* algorithms are naturally more sample efficient since they also learn from the data generated by historical policies (i.e., *experience*) (Sutton & Barto, 1998; Silver & Sutton, 2025). Following this direction, recent research has begun to incorporate off-policy data in RFT in different ways, including using nonuniform replay strategies (Li et al., 2025a), learning from positive and negative signals asymmetrically (Roux et al., 2025; Arnal et al., 2025), proposing new learning objectives based on generation consistency (Tang et al., 2025; Cohen et al., 2025), and learning from demonstrations of superior models (Yan et al., 2025), etc. Despite the efforts made by these works, off-policy RFT remains underexplored in two aspects: (1) None of these methods was compared with SOTA models on multiple mainstream math reasoning benchmarks, leaving training efficiency and final performance of these methods untested thoroughly; (2) The influence of off-policy learning on the learning process of reasoning ability remains unknown, which impedes essential understanding of off-policy learning for RFT and advancement of effective methodologies.

In this paper, we study off-policy RL for post-training finetuning of LLMs, aiming to achieve SOTA-level reasoning ability efficiently and unbox the effects of off-policy learning for useful insights. We propose **Reincarnating Mix**-policy Proximal Policy Optimization (**ReMix**), a general approach to enable on-policy proximal policy gradient methods to leverage off-policy data efficiently. ReMix consists of three major components: (1) Mix-policy proximal policy gradient with an increased Update-To-Data (UTD) ratio (Chen et al., 2021) leverages the data generated by both current policy and past policies for efficient training; (2) KL-Convex policy constraint (Ma et al., 2024b) combines the KL constraints on the base model and the precedent model to balance the trade-off between stability and flexibility during training; (3) Policy reincarnation (Agarwal et al., 2022) replaces the base model with the mix-policy RFT model in the mid way of training and restarts on-policy training, to achieve a seamless transition from efficient early-stage learning to steady asymptotic improvement. Under the synergy of the three components, ReMix is able to improve the reasoning ability of LLMs efficiently while retaining a stable and flexible training process.

In our experiments, we adopt PPO and GRPO as representative on-policy methods and implement **ReMix-PPO** and **ReMix-GRPO**. We use DeepSeek-R1-Distill-Qwen-1.5B and -7B (Guo et al., 2025) as the base models, and train our models based on DeepScaleR-Preview-Dataset (Luo et al., 2025). We conduct a range of comparative evaluations against 15 recent advanced models on five math reasoning benchmarks, including AIME’24, AMC’23, Minerva (Lewkowycz et al., 2022), OlympiadBench (He et al., 2024), and MATH500 (Hendrycks et al., 2021). Figure 1 summarizes the experimental results in a view of efficiency-performance comparison. Our method achieves an **average Pass@1 accuracy of 52.10%** (for 1.5B model) with **0.079M response rollouts, 350 training steps** and achieves **63.27%/64.39%** (for 7B model) with **0.007M/0.011M response rollouts, 50/75 training steps** respectively, showing SOTA-level performance and an over **30x to 450x training cost reduction in terms of rollout data volume**. Moreover, to gain a better understanding of off-policy learning for RFT, we conduct multifaceted studies and analysis, revealing insightful findings including the implicit preference for shorter responses due to the Whipping Effect of off-policy discrepancy, the collapse mode of self-reflection behavior under the presence of severe off-policiness, the performance under response length constraint, the impact of prompt format, etc.

108 **2 PRELIMINARIES**

109
110 **Reinforcement Learning for LLM Fine-tuning** Reinforcement Fine-Tuning (RFT) is a paradigm
111 for adapting pre-trained LLMs to specific downstream tasks using RL (Trung et al., 2024; Jaech
112 et al., 2024). In this paradigm, text generation is modeled as a Markov Decision Process (MDP)
113 $M = (\mathcal{S}, \mathcal{A}, P, R, \gamma)$, where a **state** $s_t = (q, y_{1:t}) \in \mathcal{S}$ is the prompt with the output generated so
114 far, and the **action** $a_t \in \mathcal{A}$ is the next token selected from the vocabulary \mathcal{V} . Hence, the **transition**
115 $P(s_{t+1}|s_t, a_t)$ is deterministic in this context. An episode starts from a prompt s_0 (out of a predefined
116 set \mathcal{D}_0) and terminates at an end-of-sequence token or by the maximum sequence length H .

117 The **reward** $R(s_t, a_t)$ signal is issued by either a rule-based reward function or a learned reward
118 model. In the scope of this paper, we consider the verifiable reward function. For any non-terminal
119 timestep $t < T - 1$, $R(s_t, a_t)$ is 0; on completion, the terminal reward, denoted by $R(\tau)$ for the
120 whole sequence, equals 1 if τ produces a correct and well-formatted answer and 0 otherwise. The
121 **policy** $\pi_\theta(a_t | s_t)$ in the MDP is the LLM itself, parameterized by θ , and it defines a probability
122 distribution of next-token generation. We use $d_\tau^{\pi_\theta}$ to denote the distribution of the output sequence
123 τ generated by π_θ and use $d_{s,a}^{\pi_\theta}, d_s^{\pi_\theta}$ for the state-action pairs (s, a) and the state respectively. The
124 learning objective of an RL policy is to maximize the reward function, i.e., $\pi^* = \arg \max_{\pi_\theta} J(\pi_\theta)$.

125 **Proximal Policy Gradient Methods for RFT** Proximal Policy Optimization (PPO) (Schulman
126 et al., 2017) is a canonical Policy Gradient (PG) method (Sutton & Barto, 1998) to maximize
127 $J(\pi_\theta)$, which offers stable training and implementation simplicity. PPO is further developed to be
128 GRPO (Shao et al., 2024) with a group-based advantage estimator. The policy optimization objective
129 of PPO is formulated as:

$$130 L^{\text{CLIP}}(\theta) = -\mathbb{E}_{s,a \sim d_{s,a}^{\pi_\theta}} \left[\min \left(r_\theta(s, a) \hat{A}(s, a), \text{clip}(r_\theta(s, a), 1 - \epsilon, 1 + \epsilon) \hat{A}(s, a) \right) \right], \quad (1)$$

131 where $r_\theta(s, a) = \frac{\pi_\theta(a|s)}{\pi_{\theta_{\text{old}}}(a|s)}$ represents the importance sampling ratio between the current policy π_θ
132 and the old policy $\pi_{\theta_{\text{old}}}$ (i.e., the policy before the update), $\hat{A}(s, a)$ is an estimator of the advantage
133 function with GAE (Schulman et al., 2016) as a popular choice, and the clip ratio ϵ defines the
134 clipping range that determines the proximity of policy update, thereby enhancing stability. When
135 applying RL for LLM, a KL-divergence penalty is often added to prevent the policy from deviating
136 too far from a reference model π_{base} , e.g., the SFT model. The complete objective is:

$$138 L^{\text{PPO}}(\theta) = \mathbb{E}_{s \sim d_{s,a}^{\pi_{\theta_{\text{old}}}}} \left[L^{\text{CLIP}}(\theta) + c \mathcal{H}[\pi_\theta](s) \right] + \beta \cdot \underbrace{\mathbb{E}_{s \sim d_s^{\pi_{\theta_{\text{old}}}}} [D_{\text{KL}}(\pi_\theta(\cdot | s) || \pi_{\text{base}}(\cdot | s))]}_{L_{\text{KL}}(\theta; \pi_{\text{base}})}, \quad (2)$$

139 where $\mathcal{H}[\pi_\theta](s)$ is the entropy of the policy π_θ at state s , D_{KL} is the KL metric, and c, β are weighting
140 coefficients. In this work, we view both PPO, GRPO, and other variants of PPO as Proximal Policy
141 Gradient (PPG) methods.

142 **3 REINCARNATING MIX-POLICY PROXIMAL POLICY OPTIMIZATION**

143 In this section, we introduce our method, Reincarnating Mix-policy Proximal Policy Optimization
144 (ReMix), for efficient and stable RFT of LLMs. Specifically, ReMix consists of three synergistic
145 innovations: (1) Mix-policy proximal policy gradient with an increased Update-To-Data (UTD) ratio
146 for efficient training (Section 3.1); (2) KL-Convex policy constraint to balance stability and flexibility
147 (Section 3.2); (3) Policy Reincarnation for a smooth transition from efficient early learning to stable
148 asymptotic improvement (Section 3.3). We introduce the three components below.

149 **3.1 MIX-POLICY PROXIMAL POLICY GRADIENT WITH INCREASED UTD RATIO**

150 While proximal policy gradient methods like PPO, GRPO deliver strong performance in RFT, the
151 on-policy nature of these methods leads to a significant bottleneck on data utilization. To address this
152 inefficiency, we trace back to the off-policy RL literature. To be specific, we revisit the generalized
153 proximal gradient theory (Queeney et al., 2021), which allows proximal gradient methods to make
154 use of historical trajectories generated during the past policy optimization process.

155 In this work, we launch the renaissance of off-policy RL for RFT and introduce an On-/Off-policy
156 **Mixed Proximal Policy Gradient** method (**Mix-PPG**) that strategically leverages both off-policy and

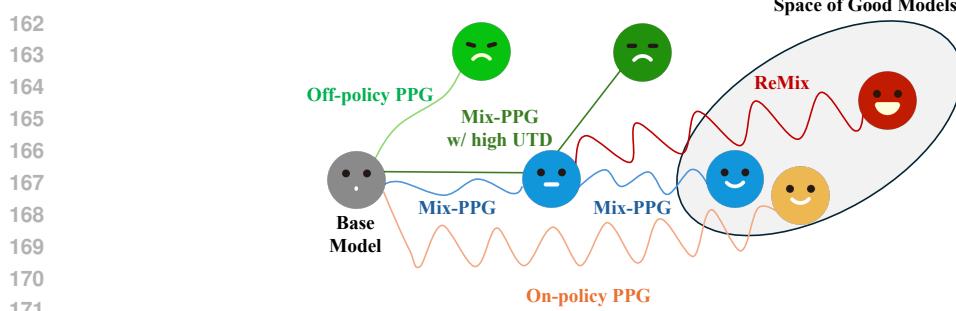


Figure 2: **The conceptual illustration of RFT for LLMs with different proximal policy gradient (PPG) methods.** Starting from a base model, (1) **on-policy PPG** methods (e.g., PPO, GRPO) train stably, yet uses data inefficiently. (2) **Off-policy PPG** is data-efficient. However, naively adopting it leads to a collapse. (3) To strike a balance, we introduce **Mix-PPG**, which boosts early-stage performance but still faces a slow asymptotic improvement and even a collapse when **adopting a high UTD ratio**. (4) Therefore, we propose policy reincarnation and introduce **ReMix**, which achieves better efficiency at no compromise of final performance.

on-policy data within a unified objective function. Formally, for policy at training step k , the mini-batch training data are sampled from a mixture of sources: the trajectories generated by historical policies (i.e., π_{k-i} for $i \sim \nu$), and the trajectories of the current policy (i.e., π_k). This hybrid sampling strategy balances two competing purposes: (1) Data Reuse: Exploiting past trajectories reduces the autoregressive rollout and inference overhead; (2) Distribution Alignment: Maintaining sufficient on-policy samples prevents training instability and degradation due to the divergence from the current state-action distribution. The policy optimization objective function can be formalized as:

$$L_k^{\text{Mix-PPG}}(\theta) = -\mathbb{E}_{i \sim \nu} \left[\mathbb{E}_{(s, a) \sim d_{s, a}^{\pi_{k-i}}} \min \left(r_{\theta}^{k-i}(s, a) A^{\pi_k}(s, a), \right. \right. \\ \left. \left. \text{clip} \left(r_{\theta}^{k-i}(s, a), \frac{\pi_k(a | s)}{\pi_{k-i}(a | s)} - \epsilon, \frac{\pi_k(a | s)}{\pi_{k-i}(a | s)} + \epsilon \right) A^{\pi_k}(s, a) \right) \right], \quad (3)$$

where $i \sim \nu$ with $i \in \{0, 1, \dots, N\}$ is a combined distribution over historical policy indices π_{k-i} and the current policy π_k (i.e., when $i = 0$), the importance sampling ratio $r_{\theta}^{k-i}(s, a) = \frac{\pi_{\theta}(a | s)}{\pi_{k-i}(a | s)}$. Notably, we incorporate a sampling strategy to strike a balance between training stability and efficient data utilization by using a portion p of off-policy data drawn from π_{k-i} and $1 - p$ on-policy data drawn from π_k with $p \in [0, 1]$. [A theoretical discussion on \$\nu\$ can be found in Appendix D.2](#). Now, we are ready to replace the on-policy policy optimization objective, e.g., the $L^{\text{Clip}}(\theta)$ term in Eq. 2, with the Mix-PPG objective $L_k^{\text{Mix-PPG}}(\theta)$ for efficient data utilization. One thing to note is, we found that explicitly maintaining the portion of on-policy data at a sufficient level is critical to effective training, as much off-policyness data will lead to a degradation or even collapse (as depicted in Figure 2).

To further improve sample efficiency, we leverage an increased Update-To-Data (UTD) ratio ([Chen et al., 2021](#)), defined originally as the number of gradient updates per environment interaction step. Specifically, we use a UTD ratio m , i.e., performing repeated gradient updates on sampled data batches for m times, thereby further reducing fresh environment interaction demands.

3.2 KL-CONVEX POLICY CONSTRAINT

Conventional RFT imposes a static KL-constraint regularization on deviations from the base pre-trained model π_{base} . This rigid static constraint fails to accommodate evolving policy distributions, which could lead to suboptimal updates during the dynamic learning process.

Inspired by the recent study ([Ma et al., 2024b](#)), we propose to dynamically update the anchor objective to a convex combination of π_{k-1} and π_{base} . On the one hand, by constraining the policy within the support of π_{base} , we enforce behavioral consistency with foundational capabilities, thereby preventing catastrophic forgetting of core skills. On the other hand, the constraint imposed on π_{k-1} serves as a dynamic adaptation to the policy's current knowledge frontier. It facilitates iterative refinement of the policy and enables the policy to continuously evolve and improve steadily. With this mechanism, the policy can leverage the strengths of both the pre-trained model and the iterative refinement process.

216 Specifically, we reconcile the KL-constraint in RFT via a KL-convex policy constraint (**KLC**),
 217 which modifies the essential optimization objective described in Eq. 2 by replacing the conventional
 218 $L_{\text{KL}}(\theta; \pi_{\text{base}})$ with the KL-convex constraint term as formulated below:
 219

$$220 L_{\text{KLC}}(\theta; \pi_{\text{base}}, k) = \mathbb{E}_s [\lambda D_{\text{KL}}(\pi_\theta(\cdot|s) \parallel \pi_{\text{base}}(\cdot|s)) + (1 - \lambda) D_{\text{KL}}(\pi_\theta(\cdot|s) \parallel \pi_{k-1}(\cdot|s))], \quad (4)$$

221 where $\lambda \in [0, 1]$ balances base-model alignment and behavioral consistency with recent policy π_{k-1} .
 222 This convex combination preserves foundational capabilities while enabling targeted adaptation, acting
 223 as a conservative regularizer against over-specialization. We provide more theoretical discussion
 224 on KLC loss in Appendix D.3. Empirically, we found that using a decaying λ consistently beat a
 225 fixed one. A linear decay schedule works well while other forms (e.g., exponential decay) did not
 226 make a significant difference.
 227

228 3.3 POLICY REINCATEINATION

230 While the mix-policy proximal PG method proposed above accelerates early-stage training, the off-
 231 policy bias in it can inevitably limit the asymptotic performance, as the empirical evidence later shown
 232 in Figure 4. Inspired by Reincarnating RL (Agarwal et al., 2022), we propose Policy Reincarnation
 233 in the context of RFT for LLMs, to seamlessly combine the advantage of both off-policy RL and
 234 on-policy RL, thus being more efficient at no cost of asymptotic performance.
 235

236 To be specific, the training process consists of the Mix-PPG stage and the reincarnating on-policy PPG
 237 stage. First, the initial policy model is trained for a predetermined T steps of gradient update according
 238 to the proposed Mix-PPG algorithm for quick improvement of policy performance. Thereafter, the
 239 reincarnation happens through two changes to the training setting: (1) reset the base model from
 240 the initial reference model π_{base} to the current policy model π_T (which alters the conventional KL
 241 constraint term), and (2) switch Mix-PPG to a on-policy PPG method (e.g., PPO or GRPO).
 242

243 Finally, by composing Mix-PPG (Eq. 3), KL-convex policy constraint (Eq. 4), and policy reincarnation,
 244 we arrive at the complete method proposed in this paper, i.e., Reincarnating Mix-policy Proximal
 245 Policy Optimization (**ReMix**), as follows:
 246

$$247 L^{\text{ReMix}}(\theta) = \begin{cases} \mathbb{E}_{d_{s,a}^{\pi_\theta}} [L^{\text{Mix-PPG}}(\theta) + c\mathcal{H}[\pi_\theta](s)] + \beta \cdot L_{\text{KLC}}(\theta; \pi_{\text{base}}, t) & \text{if } t \leq T; \\ \mathbb{E}_{d_{s,a}^{\pi_\theta}} [L^{\text{PPO}}(\theta) + c\mathcal{H}[\pi_\theta](s)] + \beta \cdot L_{\text{KLC}}(\theta; \pi_T, t) & \text{otherwise.} \end{cases} \quad (5)$$

247 Note that t is the number of batch training steps and the two changes that occur upon policy
 248 reincarnation are highlighted in blue and red respectively. In Eq. 5, we use PPO as the on-policy PPG
 249 method for demonstration. For the case of GRPO, one can replace the advantage estimation in both
 250 $L^{\text{Mix-PPG}}$ and L^{PPO} with the group-based estimation, as done in our experiments.
 251

252 The efficacy of ReMix is two-fold. First, it leverages the advantages of Mix-PPG and on-policy PPG
 253 in boosting early-stage training and stable asymptotic improvement respectively, by establishing a
 254 seamless transition between the two stages. Second, the KL-convex policy constraint and the reset
 255 of the base reference model for KL constraint (i.e., $\pi_{\text{base}} \rightarrow \pi_T$) upon policy reincarnation offers
 256 a dynamics and looser constraint compared to the conventional static KL constraint, allowing fast
 257 policy training and a larger policy optimization space. For an intuitive understanding, we provide a
 258 conceptual illustration of RFT with different proximal PG methods in Figure 2.
 259

260 4 EXPERIMENTS

261 In this section, we empirically evaluate the efficacy of ReMix on a range of commonly adopted Math
 262 reasoning benchmarks in terms of both accuracy and efficiency (Section 4.2), along with ablation
 263 studies (Section 4.3) and multifaceted analyses (Section 4.4 and Appendix L). In addition, we provide
 264 extensive evaluation regarding more evaluation metrics, code generation, base RFT algorithms, etc.
 265

266 4.1 EXPERIMENTAL SETUP

267 **Training** We use DeepSeek-R1-Distill-Qwen-1.5B and -7B (Guo et al., 2025) as the
 268 base models in our experiments. For implementation, we adopt PPO and GRPO as two representative
 269 on-policy proximal PG methods in our experiments, resulting in **ReMix-PPO** and **ReMix-GRPO**.
 270

We use DeepScaleR-Preview-Dataset (Luo et al., 2025), which comprises approximately 40,000 unique problem-answer pairs sourced from AIME (1984–2023), AMC (prior to 2023), the Omni-MATH dataset (Gao et al., 2025), and the Still dataset (Min et al., 2024). We use the DeepScaleR’s prompt format by default, instructing the LLM to follow structured step-by-step reasoning and produce a verifiable \boxed{} final answer. Full templates and instances are provided in Appendix G.

Our experiments are conducted using the `verl`¹ framework and the codebase derived from `tinyzero`². For ReMix, we use an off-policy data portion $p = 0.4$, a UTD ratio $m = 2$, a historical policy window size $N = 2$, and we set the policy reincarnation step point to $T \in \{50, 100\}$ for ReMix-PPO and $T = 50$ for ReMix-GRPO. The KL-Convex coefficient λ decays with training steps t as $\lambda(t) = \max(1 - 0.1 \cdot [\max(t - 50, 0)/10], 0.5)$. We use these configurations by default, except in hyperparameter analysis. Prompts are truncated to 766 tokens, and the maximum generation length is **8,192** tokens. The detailed hyperparameter choices are presented in Table 9.

Evaluation We evaluate the performance of different models on a series of mathematical reasoning benchmarks, including AIME’24, AMC’23, Minerva (Lewkowycz et al., 2022), OlympiadBench (He et al., 2024), and MATH500 ((Lightman et al., 2023)) (None of these datasets are contained in our train set). During evaluation, we feed the entire context into the evaluation function. The models in comparison use the same generation settings as in training, except the `do_sample` parameter is set to `false`, resulting in greedy decoding. For the evaluation of baseline methods, we use the officially released checkpoints from HuggingFace to ensure fair results; for our models, we use the best checkpoints obtained within a specific training step budget, e.g., ReMix-PPO (200 Steps).

In our experiments, we focus on the evaluation of ReMix in terms of both **model performance** and **training efficiency**. For model performance, we mainly use **Pass@1** accuracy, and Avg@32 for small datasets (AIME’24 and AMC’23) in Appendix J. For training efficiency, we evaluate the models mainly in terms of **rollout data volume**, defined as the total number of rollouts generated by the model during training, which is usually the **dominant source** of computational cost during training in practice. We also use *training steps* (i.e., the number of rollout prompt batches) and *training duration* (i.e., the actual elapsed wall-clock time) as additional aspects for efficiency evaluation.

The detailed introduction of the compared baselines is provided in Appendix F. We also provide the discussion on related off-policy methods that are infeasible to compare with in Appendix F.3. For other training details, please refer to Appendix I.

4.2 PERFORMANCE EVALUATION FOR MATH REASONING

The performance evaluation in terms of Pass@1 accuracy on five math reasoning benchmarks are shown in Table 1 and Table 2, our method ReMix achieves consistent and substantial improvements over the base 1.5B/7B model on all five benchmarks. For **ReMix-PPO**, it achieves an **average performance gain of 14.52 points and 12.31 points** over 1.5B and 7B base models respectively, **achieving the second-best average score for 1.5B and the best for 7B among all the baselines**. In addition, compared with PPO (900 Steps, 1.5B) and PPO (200 Steps, 7B), our model achieves higher average scores within 100 steps for 1.5B and 50 steps for 7B. Similarly, our model exceeds GRPO (100 Steps, 1.5B) and GRPO (200 Steps, 1.5B) within 50 and 200 training steps, respectively. This indicates that ReMix is able to achieve competitive reasoning ability efficiently with overall no compromise in accuracy and even showing a higher accuracy.

More importantly, we move on to the evaluation in terms of training efficiency. This is shown in the last volume (i.e., Cost) of Table 1 and 2, and notably, Figure 1 illustrates the **efficiency–accuracy trade-off** in terms of **rollout data volume (log₁₀ scale)** versus **average Pass@1 accuracy**, where the scores are out of Table 1 and 2 (i.e., **Avg.** and **Cost**). In the ideal case, the model should appear in the top-left corner of the plot. To ensure a fair comparison, the rollout data volume of square-marked models (which means the models fine-tuned upon DeepScaleR) includes the data cost of training DeepScaleR itself. For ReMix-GRPO and GRPO, we report results after 200 training steps due to computational resource constraints.

Specifically, for 1.5B models, **ReMix-PPO** matches **DeepScaleR**, the strongest competitor, with just 0.079M vs. 2.519M rollouts, over a **30x reduction** in rollout data volume. Also, **ReMix-PPO**

¹<https://github.com/volcengine/verl>

²<https://github.com/Jiayi-Pan/TinyZero>

324
 325
 326
 327
**Table 1: Pass@1 accuracy (%) and training cost (in terms of Rollout Data Volume) of 1.5B
 328 models. Bolded and underlined values denote the highest and the second-highest scores in each
 329 dataset (i.e., column). ‘-’ denotes that not enough information was found. ReMix achieves better
 330 average scores than both the standard PPO and GRPO in a significantly more efficient manner.**

Model	AIME’24	AMC’23	MATH500	Minerva	Olympiad	Avg.	Cost
R1-Distill-Qwen-1.5B (Base)	33.33	43.37	67.40	16.54	27.26	37.58	N/A
Open-RS1	23.33	42.17	64.20	16.18	27.11	34.60	0.058M
Open-RS2	16.67	45.78	65.00	18.38	26.96	34.56	0.029M
Open-RS3	16.67	44.58	67.60	15.64	25.48	33.99	0.029M
AdaptThink	13.33	57.83	78.60	23.90	38.07	42.35	0.643M
II-Thought	26.67	56.63	73.00	23.16	40.89	44.07	-
FASTCuRL-preview	26.67	60.24	74.20	20.22	32.59	42.78	0.676M
FASTCuRL-V3	36.67	66.27	84.40	28.67	43.56	51.91	2.478M
L1-Exact*	23.33	71.08	84.00	<u>29.41</u>	<u>44.59</u>	50.48	3.953M
L1-Max*	20.00	<u>69.88</u>	83.00	29.04	46.37	49.66	2.764M
DeepScaleR	40.00	65.06	83.20	29.04	43.41	52.14	2.519M
GRPO (100 Steps)	30.00	56.63	75.80	25.37	38.22	45.20	0.205M
GRPO (200 Steps)	36.67	61.45	80.00	25.37	39.70	48.64	0.410M
ReMix-GRPO (50 Steps)	23.33	57.83	80.40	26.10	39.70	45.47	0.061M
ReMix-GRPO (100 Steps)	23.33	62.65	82.00	28.68	39.70	47.27	0.163M
ReMix-GRPO (200 Steps)	33.33	65.06	84.60	26.10	43.55	50.53	0.368M
PPO (500 Steps)	36.67	62.65	82.60	25.73	40.14	49.56	0.128M
PPO (900 Steps)	30.00	<u>69.88</u>	84.00	25.74	43.41	50.61	0.230M
ReMix-PPO (100 Steps)	<u>43.33</u>	63.86	79.60	26.84	39.41	50.61	0.020M
ReMix-PPO (200 Steps)	46.67	62.65	82.20	26.10	40.74	51.67	0.041M
ReMix-PPO (350 Steps)	36.67 ^[3.34]	69.88 ^{↑26.51}	82.00 ^{↑14.60}	30.15 ^[13.61]	41.78 ^{↑14.52}	52.10 ^{↑14.52}	0.079M

345
 346
Table 2: Pass@1 accuracy (%) and training cost (in terms of Rollout Data Volume) of 7B models.
347 ReMix-PPO achieves the best average score within 75 training steps.

Model	AIME’24	AMC’23	MATH500	Minerva	Olympiad	Avg.	Cost
R1-Distill-Qwen-7B (Base)	33.33	68.68	83.80	30.15	44.44	52.08	N/A
ReasonFlux-F1	20.00	54.22	77.20	29.04	37.04	43.50	-
Light-R1	30.00	66.27	87.00	34.56	47.56	53.08	-
Skywork-OR1-Preview	43.33	63.86	84.40	29.41	46.22	53.44	>8.192M
Polaris	40.00	63.86	87.60	36.40	48.00	55.17	-
AdaptThink	46.67	75.90	87.60	33.46	50.22	58.77	0.307M
AceReason-Nemotron	60.00	80.72	89.00	36.40	50.07	63.24	>3.584M
ReMix-GRPO (75 Steps)	63.88	90.60	80.72	40.07	53.78	65.81	0.046M
ReMix-GRPO (200 Steps)	64.37	91.60	81.93	39.34	53.19	66.09	0.163M
PPO (50 Steps)	33.33	71.08	87.20	36.03	48.00	55.13	0.013M
PPO (100 Steps)	40.00	77.11	<u>90.00</u>	35.66	<u>51.56</u>	58.87	0.026M
PPO (200 Steps)	53.33	78.31	87.00	34.19	48.88	60.34	0.051M
ReMix-PPO (25 Steps)	36.67	78.31	89.00	<u>38.24</u>	50.22	58.49	0.003M
ReMix-PPO (50 Steps)	56.66	<u>79.52</u>	88.60	38.97	52.59	63.27	0.007M
ReMix-PPO (75 Steps)	63.33 ^[30.00]	78.31 ^[9.63]	90.20 ^[7.40]	37.50 ^{↑7.35}	52.59 ^{↑8.15}	64.39 ^{↑12.31}	0.011M

361 reaches 50.61 after 0.020M rollouts, **10x fewer** than PPO (50.61 at 0.230M), highlighting rapid early
 362 gains. For 7B models, **ReMix-PPO** topping AceReason-Nemotron with over a **450x reduction**
 363 in rollout data volume, and outperforming PPO with a **6x reduction**. Notably, the average rollout
 364 response length of ReMix is lower than the baseline models (see Fig. 4), hence the exact efficiency
 365 should be higher. The corresponding detailed factors associated with computational cost for training
 366 all compared models above are shown in Appendix I.

367 **Takeaway 1. ReMix can learn strong reasoning ability in a highly efficient way.**

368
 369 ReMix achieves SOTA-level accuracies at 1.5B and 7B scales on five math reasoning benchmarks,
 370 with an over 6x to 10x reduction in rollout data volume when outperforming PPO and
 371 an over 30x to 450x reduction when performing on par with (or exceeding) the best baseline.

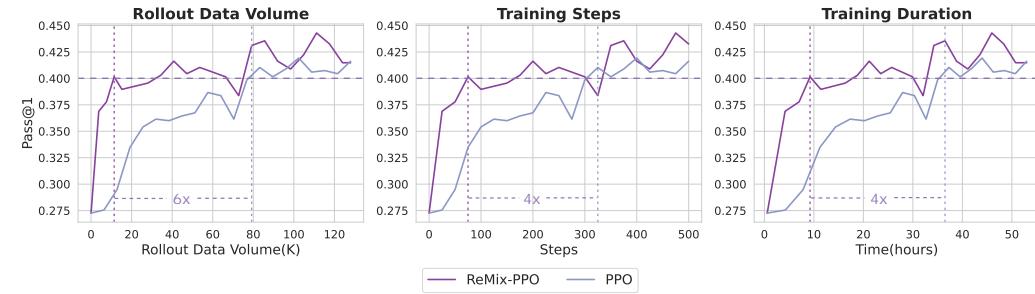
372
 373 **4.3 ABLATION STUDIES**

374 To assess the contribution of each components in ReMix, we conduct ablation studies focusing on
 375 both training dynamics and final performance. We use ReMix-PPO for the ablation studies.
 376

377 The results of the ablation studies regarding Pass@1 accuracy are presented in Table 3. First, when
 378 Mix-PPG, the core of ReMix, is ablated, the method degenerates to PPO since it does not make sense

378 **Table 3: Ablation studies regarding Pass@1 accuracy.** The three components of ReMix work in
 379 synergy for both efficiency and final performance.

380 Model	381 AIME’24	382 AMC’23	383 MATH500	384 Minerva	385 Olympiad	386 Avg.
381 R1-Distill-Qwen-1.5B (Base Model)	381 33.33	381 43.37	381 67.40	381 16.54	381 27.26	381 37.58
382 PPO (500 Steps)	382 36.67	382 62.65	382 82.60	382 25.73	382 40.14	382 49.56
383 ReMix-PPO (350 Steps)	383 36.67	383 69.88	383 82.00	383 30.15	383 41.78	383 52.10
384 ReMix-PPO w/o UTD	384 36.67	384 62.65	384 82.20	384 28.68	384 42.96	384 50.63
385 ReMix-PPO w/o KL-Convex	385 30.00	385 65.06	385 81.60	385 27.94	385 42.22	385 49.36
386 ReMix-PPO w/o Policy Reincarnation	386 20.00	386 67.47	386 82.00	386 26.84	386 40.00	386 47.26
386 ReMix-PPO w/o UTD, KL-Convex, Policy Reincarnation	386 40.00	386 57.83	386 80.40	386 25.74	386 39.55	386 48.70



387 **Figure 3: Training efficiency comparison for ReMix-PPO and PPO (1.5B) on Olympiad.** We
 388 evaluate training efficiency across three dimensions: rollout data volume, training steps, and training
 389 duration. ReMix achieves a score above 40%, around **4x to 6x faster** than PPO.

390 any longer to apply other components of ReMix. Removing any single component among increased
 391 UTD, KL-Convex, and policy reincarnation leads to a final average score comparable to PPO but
 392 lower than ReMix within 500 training steps. Dropping policy reincarnation hurt the most among the
 393 three ablations. Further, when Mix-PPG works solely, it leads to an even lower score.

394 This reflects the off-policy nature in Mix-PPG: although it significantly increases the training effi-
 395 ciency, the off-policyness bias may hinder the convergence performance. The superiority in efficiency
 396 brought by Mix-PPG can be observed by referring to the first subplot of Figure 4: Mix-PPG shows a
 397 somewhat surprising boost of Pass@1 accuracy within the first 100 training steps, which an increased
 398 UTD further enhances it; while the KL-Convex and policy reincarnation in ReMix contribute to the
 399 steady asymptotic improvement.

400 **Takeaway 2. The three components work in synergy for efficiency and final performance.**

401 Mix-PPG with an increased UTD boosts early-stage training significantly, while policy
 402 reincarnation plays a critical role to ensure asymptotic improvement.

403 **Training Curves** In addition to the efficiency evaluation in terms of rollout data volume, we present
 404 the training curves for ReMix-PPO and PPO in Figure 3 on Olympiad regarding two more efficiency
 405 aspects, i.e., training steps and wall-clock time. Our method demonstrates superior training efficiency
 406 by achieving a score above 40 on Olympiad with a 6x and 4x reduction in rollout data volume and
 407 wall-clock time. We provide more training curves for the other four benchmarks in Appendix K.

408 **4.4 ANALYSIS ON THE EFFECTS OF OFF-POLICY RL FOR LLM RFT**

409 In this subsection, we present an empirical analysis to gain better understanding of the effects of
 410 off-policy RL on LLM RFT. For convenience, we use ReMix-PPO for the analysis in the following.

411 To delve into the influence of off-policy RL enabled by ReMix on the reasoning behaviors during the
 412 learning process of LLMs, we make use of two more metrics: relative response length (against the
 413 training dynamics of PPO), and **self-reflection rate** that is calculated according to the occurrence of
 414 reflection tokens (e.g., ‘verify’, ‘check’, ‘but’, ‘wait’, etc.). Moreover, we compare PPO, ReMix-PPO,
 415 Mix-PPG and Mix-PPG with an increased UTD. The results are shown in Figure 4.

416 The vanilla PPO shows a steady increase of Pass@1 accuracy as well as a decrease in response
 417 length, while maintaining a self-reflection rate near 1. Mix-PPG accelerates early training but yields
 418 inferior asymptotic performance (see more in Figure 9), with a clear drop in response length and
 419 self-reflection rate. When applying an increased UTD ratio, Mix-PPG speeds further yet shows a

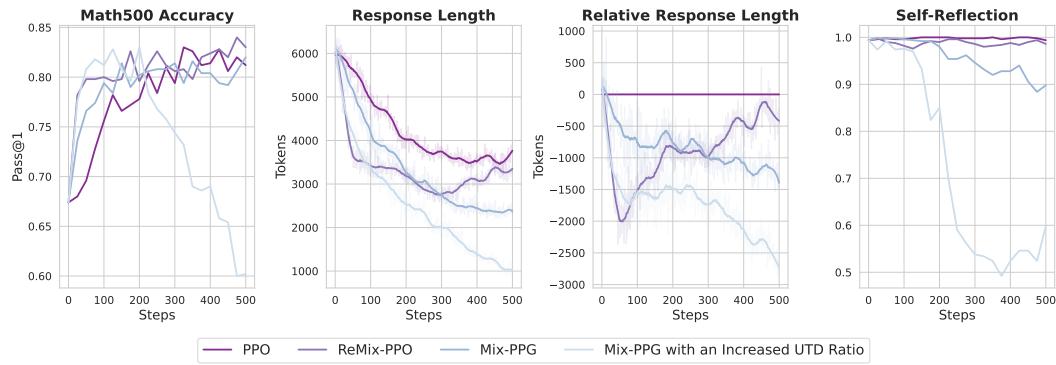


Figure 4: **Training dynamics regarding accuracy, response length, self-reflection rate for on-policy v.s. off-policy training.** ReMix shows a merged learning behavior and **perfectly combines the superior efficiency and the asymptotic improvement** thanks to the policy reincarnation.

destructive degradation after 200 steps, accompanied by a sharp decrease in response length and self-reflection rate. This pathology drives the model to generate a final answer without adequate intermediate deliberation, resulting in a drop of accuracy. Please see Appendix N for detailed cases.

ReMix seems to perfectly combine the early-stage efficiency of Mix-PPG with an increased UTD and the asymptotic improvement of PPO, thanks to policy reincarnation. Throughout training, ReMix first quickly decreases the response length and improve its accuracy in the early stage; it then lengthens responses and uses more reflection for careful exploration and further improvement of accuracy.

Takeaway 3. A trade-off between efficiency and final performance needs to be balanced when incorporating off-policy training in RFT.

More off-policy training leads to a faster early-stage boost with a larger policy shift, leading to shorter responses and quicker unlearning of self-reflection, consequently damaging reasoning performance. ReMix well leverages off-policy efficiency at no compromise of final accuracy.

The Implicit Preference of Off-policy Learning for Shorter Responses To take a further step on why off-policy learning leads to the observed reasoning behaviors, we conduct a formal analysis on the learning dynamics when optimizing the Mix-PPG loss function $L_k^{\text{Mix-PPG}}(\theta)$ (shown in Eq. 3). Similarly as in (Fatemi et al., 2025), the average loss of Mix-PPG can be formulated below:

$$L_{\text{Avg}}^{\text{Mix-PPG}} = \frac{1}{H} \sum_{h=0}^H L_h^{\text{Mix-PPG}} \propto -\frac{1}{H} \sum_{h=0}^H r_{\theta}^{k-i} A_h^{\pi_k} \quad (6)$$

A simple derivation is provided in Appendix D.4. With the equation above, we can find: when the advantage estimate is *negative*, the model learns to minimize the loss by steering its policy to achieve a lower importance sampling ratio. As the policy loss is almost always positive (as in Figure 10), the advantage estimates are negative most of the time in our experiments. Also, the importance sampling ratio stays above one empirically (as in Figure 9), directly amplifies the loss term. Since the average loss is computed based on the data of historical policy π_{k-i} , there apparently exists a *Whipping Effect*: the longer the response is, the larger the distribution shift should be on later states. Consequently, the model tends to prefer shorter responses to reduce the loss associated with long rollout trajectories. This tendency is further amplified as the proportion of off-policy data increases.

4.5 MORE RESULTS ON PASS@K METRICS, CODE GENERATION, OTHER BASE ALGORITHMS

Pass@K Evaluation In the experiments above, we mainly use **Pass@1** accuracy (as well as **Avg@32**) as the evaluation metric. As advocated in (Yue et al., 2025), we extend our evaluation by

Table 4: **Pass@8 / Pass@16 accuracy (%) of 1.5B models on five reasoning benchmarks.**

Model	AIME'24	AMC'23	MATH500	Minerva	Olympiad	Avg.
R1-Distill-Qwen-1.5B (Base, Pass@8)	20.00	44.58	73.00	19.12	31.70	37.68
ReMix-PPO (Pass@8)	30.00	68.67	84.60	28.68	46.67	51.72
R1-Distill-Qwen-1.5B (Base, Pass@16)	20.00	48.19	75.40	18.75	32.59	38.99
ReMix-PPO (Pass@16)	30.00	72.29	86.80	30.15	46.96	53.24

486 **Table 5: Pass@8 / Pass@16 accuracy (%) of 7B models on five reasoning benchmarks.**

Model	AIME'24	AMC'23	MATH500	Minerva	Olympiad	Avg.
R1-Distill-Qwen-7B (Base, Pass@8)	63.33	83.13	93.80	47.79	55.11	68.63
ReMix-PPO (Pass@8)	76.67	92.77	95.80	50.74	63.70	75.94
R1-Distill-Qwen-7B (Base, Pass@16)	73.33	85.54	94.80	51.10	57.19	72.39
ReMix-PPO (Pass@16)	80.00	92.77	95.40	53.31	65.33	77.36

492 **Table 6: Performance evaluation of ReMix based on Dr.GRPO.** Pass@1 accuracy (%) on five
493 benchmarks (Avg@32 for AIME'24/AMC'23). The max response length is 4096.

Model	AIME'24	AMC'23	MATH500	Minerva	Olympiad	Avg.	Cost
R1-Distill-Qwen-1.5B (Base Model)	9.38	32.42	63.40	18.75	22.67	29.32	N/A
Dr.GRPO (350 steps)	21.25	63.10	79.40	27.21	41.33	46.48	0.358 M
Dr.GRPO (400 steps)	26.25	59.94	80.60	26.57	41.03	46.86	0.409 M
ReMix-Dr.GRPO (325 steps)	26.25	63.25	81.60	27.21	41.63	47.09	0.291 M
ReMix-Dr.GRPO (400 steps)	28.75	62.65	82.00	27.74	42.96	48.42	0.368 M

500 using **Pass@8/16** accuracy here. The results are summarized in Table 4 and 5. We can observe that
501 ReMix-PPO effectively improves both the Pass@8 and Pass@16 accuracies of the base model.
502

503 **ReMix for Dr.GRPO** To further examine the generalization of ReMix regarding different base
504 algorithms, we evaluate the effect of ReMix based on Dr.GRPO (Liu et al., 2025b). For the convenience
505 of experimentation, we use DeepSeek-R1-Distill-Qwen-1.5B as the base model and set the max response length to 4096. We evaluate the models with sampling temperature of 0.7. All other experimental settings follow those in Section 4.1. The results show that ReMix-Dr.GRPO
506 delivers consistent improvements over Dr.GRPO with higher efficiency, mirroring the finding in our
507 main evaluation above. This demonstrates the generality of ReMix as an effective drop-in method.
508

509 **ReMix for Code Generation** Beyond Math reasoning, we move on to code generation to evaluate the domain generalization of ReMix. We use Skywork-OR1-RL-Data (He et al., 2025) for training. We set the max prompt length to 1600 because coding prompts are longer than math tasks. We evaluate our models on LiveCodeBench(8/1/24-2/1/25) (Jain et al., 2025) with sampling temperature of 0.7. All other experimental settings follow those in Section 4.1. The results are shown in Table 7. For both 1.5B and 7B scales, ReMix-PPO improves PPO while using much less cost, which aligns with our observation in the Math reasoning domain. This demonstrates the generalization ability of our method across both domains.

523 **ReMix for Llama-series Base Model** In addition to the Qwen-series base model used above, we
524 also evaluate the effect of ReMix based on DeepSeek-R1-Distill-Llama-8B in Table 16 in the appendix. Similarly, we found that ReMix-PPO improves the performance of the base model and
525 outperforms PPO in terms of both score and efficiency.

527 **Other Analysis** We provide more analysis on the performance under response length constraint,
528 the impact of prompt format, etc. Please refer to Appendix L for complete analysis results.

5 CONCLUSION

532 In this paper, we aim to address the notorious drawback of on-policy RFT methods (e.g., PPO and
533 GRPO) on training inefficiency and prohibitive computational cost. We launch the renaissance of
534 off-policy RL and propose Reincarnating Mix-policy Proximal Policy Gradient (ReMix), a general
535 approach to enable on-policy RFT methods like PPO and GRPO to leverage off-policy data. In our
536 experiments, we implement ReMix upon PPO, GRPO, and 1.5B-, 7B-scale base models. Through
537 evaluating the reasoning accuracy and training efficiency of ReMix on five math reasoning benchmarks
538 against 15 recent advanced baseline models, we demonstrate the superiority of ReMix in improving
539 training efficiency and achieving SOTA-level reasoning performance with a great reduction in training
cost. Due to the space constraint, we provide the discussion on limitations in Appendix C.

540 ETHICS STATEMENT
541542 We adhere to the code of ethics and the general principles. Our study fine-tunes publicly released
543 LLM using reinforcement learning to improve mathematical reasoning, which involves no human
544 subjects, user studies, or personal data. All datasets used for training and evaluation are publicly
545 available. These resources contain no personally identifiable information or otherwise harmful
546 information.547 To mitigate potential risk, our release will be research-only. The authors declare no conflicts of
548 interest and no sponsorship that would unduly influence the research.
549550 REPRODUCIBILITY STATEMENT
551552 For reproducibility, we release an anonymous repository ([https://anonymous.4open.
553 science/r/anonymous-remix-2025](https://anonymous.4open.science/r/anonymous-remix-2025)) containing evaluation pipelines and our trained mod-
554 els. All experimental settings, hyperparameters and datasets are listed in Subsection 4.1 and Ap-
555 pendix I; descriptions of compared baselines can be seen in Appendix F. All datasets, base models
556 and baselines we used are publicly available in HuggingFace.
557558 REFERENCES
559560 Rishabh Agarwal, Max Schwarzer, Pablo Samuel Castro, Aaron C. Courville, and Marc G. Bellemare.
561 Reincarnating reinforcement learning: Reusing prior computation to accelerate progress. In
562 *NeurIPS*, 2022.563 Pranjal Aggarwal and Sean Welleck. L1: Controlling how long a reasoning model thinks with
564 reinforcement learning. *arXiv preprint arXiv:2503.04697*, 2025.
565566 Arash Ahmadian, Chris Cremer, Matthias Gallé, Marzieh Fadaee, Julia Kreutzer, Olivier Pietquin,
567 Ahmet Üstün, and Sara Hooker. Back to basics: Revisiting reinforce-style optimization for learning
568 from human feedback in llms. In *ACL*, pp. 12248–12267, 2024.569 Chenxin An, Zhihui Xie, Xiaonan Li, Lei Li, Jun Zhang, Shansan Gong, Ming Zhong, Jingjing
570 Xu, Xipeng Qiu, Mingxuan Wang, and Lingpeng Kong. Polaris: A post-training recipe for
571 scaling reinforcement learning on advanced reasoning models, 2025. URL <https://hunlp.github.io/blog/2025/Polaris>.
572573 Charles Arnal, GaĂĂtan Naroziak, Vivien Cabannes, Yunhao Tang, Julia Kempe, and Remi Munos.
574 Asymmetric reinforce for off-policy reinforcement learning: Balancing positive and negative
575 rewards. *arXiv preprint arXiv:2506.20520*, 2025.
576577 Yuntao Bai, Andy Jones, Kamal Ndousse, Amanda Askell, Anna Chen, Nova DasSarma, Dawn Drain,
578 Stanislav Fort, Deep Ganguli, Tom Henighan, et al. Training a helpful and harmless assistant with
579 reinforcement learning from human feedback. *arXiv preprint arXiv:2204.05862*, 2022a.580 Yuntao Bai, Saurav Kadavath, Sandipan Kundu, Amanda Askell, Jackson Kernion, Andy Jones, Anna
581 Chen, Anna Goldie, Azalia Mirhoseini, Cameron McKinnon, et al. Constitutional ai: Harmlessness
582 from ai feedback. *arXiv preprint arXiv:2212.08073*, 2022b.
583584 Xinyue Chen, Che Wang, Zijian Zhou, and Keith W Ross. Randomized ensembled double q-learning:
585 Learning fast without a model. In *ICLR*, 2021.586 Yang Chen, Zhuolin Yang, Zihan Liu, Chankyu Lee, Peng Xu, Mohammad Shoeybi, Bryan Catanzaro,
587 and Wei Ping. Acereason-nemotron: Advancing math and code reasoning through reinforcement
588 learning. *arXiv preprint arXiv:2505.16400*, 2025.589 Taco Cohen, David W Zhang, Kunhao Zheng, Yunhao Tang, Remi Munos, and Gabriel Synnaeve. Soft
590 policy optimization: Online off-policy rl for sequence models. *arXiv preprint arXiv:2503.05453*,
591 2025.
592593 Quy-Anh Dang and Chris Ngo. Reinforcement learning for reasoning in small llms: What works and
what doesn't. *arXiv preprint arXiv:2503.16219*, 2025.

594 Lasse Espeholt, Hubert Soyer, Rémi Munos, Karen Simonyan, Volodymyr Mnih, Tom Ward, Yotam
 595 Doron, Vlad Firoiu, Tim Harley, Iain Dunning, Shane Legg, and Koray Kavukcuoglu. IMPALA:
 596 scalable distributed deep-rl with importance weighted actor-learner architectures. In *ICML*,
 597 volume 80, pp. 1406–1415, 2018.

598 Mehdi Fatemi, Banafsheh Rafiee, Mingjie Tang, and Kartik Talamadupula. Concise reasoning via
 599 reinforcement learning. *arXiv preprint arXiv:2504.05185*, 2025.

601 S. Fujimoto, H. v. Hoof, and D. Meger. Addressing function approximation error in actor-critic
 602 methods. In *ICML*, volume 80, pp. 1582–1591, 2018.

604 Bofei Gao, Feifan Song, Zhe Yang, Zefan Cai, Yibo Miao, Qingxiu Dong, Lei Li, Chenghao Ma,
 605 Liang Chen, Runxin Xu, Zhengyang Tang, Benyou Wang, Daoguang Zan, Shanghaoran Quan,
 606 Ge Zhang, Lei Sha, Yichang Zhang, Xuancheng Ren, Tianyu Liu, and Baobao Chang. Omni-math:
 607 A universal olympiad level mathematic benchmark for large language models. In *ICLR*, 2025.

608 Daya Guo, Dejian Yang, Huawei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu,
 609 Shirong Ma, Peiyi Wang, Xiao Bi, et al. Deepseek-r1: Incentivizing reasoning capability in llms
 610 via reinforcement learning. *arXiv preprint arXiv:2501.12948*, 2025.

611 Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. Soft actor-critic: Off-policy
 612 maximum entropy deep reinforcement learning with a stochastic actor. In *ICML*, pp. 1861–1870.
 613 Pmlr, 2018.

615 Chaoqun He, Renjie Luo, Yuzhuo Bai, Shengding Hu, Zhen Thai, Junhao Shen, Jinyi Hu, Xu Han,
 616 Yujie Huang, Yuxiang Zhang, Jie Liu, Lei Qi, Zhiyuan Liu, and Maosong Sun. OlympiadBench: A
 617 challenging benchmark for promoting AGI with olympiad-level bilingual multimodal scientific
 618 problems. In *ACL*, pp. 3828–3850, 2024.

619 Jujie He, Jiacai Liu, Chris Yuhao Liu, Rui Yan, Chaojie Wang, Peng Cheng, Xiaoyu Zhang, Fuxiang
 620 Zhang, Jiacheng Xu, Wei Shen, Siyuan Li, Liang Zeng, Tianwen Wei, Cheng Cheng, Bo An, Yang
 621 Liu, and Yahui Zhou. Skywork open reasoner 1 technical report. *arXiv preprint arXiv:2505.22312*,
 622 2025.

623 Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song,
 624 and Jacob Steinhardt. Measuring mathematical problem solving with the MATH dataset. In
 625 *NeurIPS*, 2021.

627 Matteo Hessel, Joseph Modayil, Hado van Hasselt, Tom Schaul, Georg Ostrovski, Will Dabney,
 628 Dan Horgan, Bilal Piot, Mohammad Gheshlaghi Azar, and David Silver. Rainbow: Combining
 629 improvements in deep reinforcement learning. In Sheila A. McIlraith and Kilian Q. Weinberger
 630 (eds.), *AAAI*, pp. 3215–3222, 2018.

631 Takuwa Hiraoka, Takahisa Imagawa, Taisei Hashimoto, Takashi Onishi, and Yoshimasa Tsuruoka.
 632 Dropout q-functions for doubly efficient reinforcement learning. *arXiv preprint arXiv:2110.02034*,
 633 2021.

635 Intelligent-Internet. Ii-thought. <https://ii.inc/web/blog/post/ii-thought>, 2025.

636 Aaron Jaech, Adam Kalai, Adam Lerer, Adam Richardson, Ahmed El-Kishky, Aiden Low, Alec
 637 Helyar, Aleksander Madry, Alex Beutel, Alex Carney, et al. Openai o1 system card. *arXiv preprint*
 638 *arXiv:2412.16720*, 2024.

640 Naman Jain, King Han, Alex Gu, Wen-Ding Li, Fanjia Yan, Tianjun Zhang, Sida Wang, Armando
 641 Solar-Lezama, Koushik Sen, and Ion Stoica. Livecodebench: Holistic and contamination free
 642 evaluation of large language models for code. In *ICLR*, 2025.

643 Ziwei Ji, Tiezheng Yu, Yan Xu, Nayeon Lee, Etsuko Ishii, and Pascale Fung. Towards mitigating
 644 hallucination in large language models via self-reflection. *arXiv preprint arXiv:2310.06271*, 2023.

645 Kimi, Angang Du, Bofei Gao, Bowei Xing, Changjiu Jiang, Cheng Chen, Cheng Li, Chenjun Xiao,
 646 Chenzhuang Du, Chonghua Liao, et al. Kimi k1. 5: Scaling reinforcement learning with llms.
 647 *arXiv preprint arXiv:2501.12599*, 2025.

648 Aitor Lewkowycz, Anders Andreassen, David Dohan, Ethan Dyer, Henryk Michalewski, Vinay Ra-
 649 masesh, Ambrose Slone, Cem Anil, Imanol Schlag, Theo Gutman-Solo, et al. Solving quantitative
 650 reasoning problems with language models. *NeurIPS*, 35:3843–3857, 2022.

651

652 Siheng Li, Zhanhui Zhou, Wai Lam, Chao Yang, and Chaochao Lu. Repo: Replay-enhanced policy
 653 optimization. *arXiv preprint arXiv:2506.09340*, 2025a.

654

655 Zhong-Zhi Li, Duzhen Zhang, Ming-Liang Zhang, Jiaxin Zhang, Zengyan Liu, Yuxuan Yao, Haotian
 656 Xu, Junhao Zheng, Pei-Jie Wang, Xiuyi Chen, et al. From system 1 to system 2: A survey of
 657 reasoning large language models. *arXiv preprint arXiv:2502.17419*, 2025b.

658

659 Hunter Lightman, Vineet Kosaraju, Yura Burda, Harri Edwards, Bowen Baker, Teddy Lee, Jan
 660 Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. Let’s verify step by step. *arXiv preprint*
 661 *arXiv:2305.20050*, 2023.

662

663 Jinyi Liu, Yifu Yuan, Jianye Hao, Fei Ni, Lingzhi Fu, Yibin Chen, and Yan Zheng. Enhancing
 664 robotic manipulation with ai feedback from multimodal large language models. *arXiv preprint*
 665 *arXiv:2402.14245*, 2024.

666

667 Jinyi Liu, Yan Zheng, Rong Cheng, Qiyu Wu, Wei Guo, Fei Ni, Hebin Liang, Yifu Yuan, Hangyu
 668 Mao, Fuzheng Zhang, et al. From chaos to order: The atomic reasoner framework for fine-grained
 669 reasoning in large language models. *arXiv preprint arXiv:2503.15944*, 2025a.

670

671 Zichen Liu, Changyu Chen, Wenjun Li, Penghui Qi, Tianyu Pang, Chao Du, Wee Sun Lee, and Min
 672 Lin. Understanding r1-zero-like training: A critical perspective. *arXiv preprint arXiv:2503.20783*,
 673 2025b.

674

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

678

679 Yi Ma, Hongyao Tang, Dong Li, and Zhaopeng Meng. Reining generalization in offline reinforcement
 680 learning via representation distinction. *NeurIPS*, 36:40773–40785, 2023.

681

682 Yi Ma, Jianye Hao, Xiaohan Hu, Yan Zheng, and Chenjun Xiao. Iteratively refined behavior
 683 regularization for offline reinforcement learning. In *NeurIPS*, 2024a.

684

685 Yi Ma, Jianye Hao, Xiaohan Hu, Yan Zheng, and Chenjun Xiao. Iteratively refined behavior
 686 regularization for offline reinforcement learning. *NeurIPS*, 37:56215–56243, 2024b.

687

688 Yingqian Min, Zhipeng Chen, Jinhao Jiang, Jie Chen, Jia Deng, Yiwen Hu, Yiru Tang, Jiapeng Wang,
 689 Xiaoxue Cheng, Huatong Song, Wayne Xin Zhao, Zheng Liu, Zhongyuan Wang, and Ji-Rong Wen.
 690 Imitate, explore, and self-improve: A reproduction report on slow-thinking reasoning systems,
 691 2024.

692

693 OpenAI. Gpt-3.5. Technical report, OpenAI, 2022. URL <https://platform.openai.com/docs/models/gpt-3-5>.

694

695 James Queeney, Yannis Paschalidis, and Christos G. Cassandras. Generalized proximal policy
 696 optimization with sample reuse. In *NeurIPS*, pp. 11909–11919, 2021.

697

698 Carlo Romeo, Girolamo Macaluso, Alessandro Sestini, and Andrew D Bagdanov. Speq: Offline
 699 stabilization phases for efficient q-learning in high update-to-data ratio reinforcement learning. In
 700 *RLC*, 2021.

701

Nicolas Le Roux, Marc G Bellemare, Jonathan Lebensonold, Arnaud Bergeron, Joshua Greaves,
 702 Alex Fréchette, Carolyne Pelletier, Eric Thibodeau-Laufer, Sándor Toth, and Sam Work. Ta-
 703 pered off-policy reinforce: Stable and efficient reinforcement learning for llms. *arXiv preprint*
 704 *arXiv:2503.14286*, 2025.

705

John Schulman, Sergey Levine, Pieter Abbeel, Michael I. Jordan, and Philipp Moritz. Trust region
 706 policy optimization. In *ICML*, volume 37, pp. 1889–1897, 2015.

702 John Schulman, Philipp Moritz, Sergey Levine, Michael I. Jordan, and Pieter Abbeel. High-
 703 dimensional continuous control using generalized advantage estimation. In *ICLR*, 2016.

704

705 John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy
 706 optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.

707 ByteDance Seed, Jiaze Chen, Tiantian Fan, Xin Liu, Lingjun Liu, Zhiqi Lin, Mingxuan Wang,
 708 Chengyi Wang, Xiangpeng Wei, Wenyuan Xu, et al. Seed1. 5-thinking: Advancing superb
 709 reasoning models with reinforcement learning. *arXiv preprint arXiv:2504.13914*, 2025.

710

711 Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Mingchuan Zhang, Y. K. Li,
 712 Y. Wu, and Daya Guo. Deepseekmath: Pushing the limits of mathematical reasoning in open
 713 language models. *arXiv preprint arXiv:2402.03300*, 2024.

714 David Silver and Richard S. Sutton. Welcome to the era of experience, 2025.

715

716 Mingyang Song, Mao Zheng, Zheng Li, Wenjie Yang, Xuan Luo, Yue Pan, and Feng Zhang. Fastcurl:
 717 Curriculum reinforcement learning with progressive context extension for efficient training r1-like
 718 reasoning models. *arXiv preprint arXiv:2503.17287*, 2025.

719 Richard S. Sutton and Andrew G. Barto. *Reinforcement learning - an introduction*. Adaptive
 720 computation and machine learning. MIT Press, 1998. ISBN 978-0-262-19398-6.

721 Yunhao Tang, Taco Cohen, David W Zhang, Michal Valko, and Rémi Munos. Rl-finetuning llms
 722 from on-and off-policy data with a single algorithm. *arXiv preprint arXiv:2503.19612*, 2025.

723

724 Luong Trung, Xinbo Zhang, Zhanming Jie, Peng Sun, Xiaoran Jin, and Hang Li. Reft: Reasoning
 725 with reinforced fine-tuning. In *ACL*, pp. 7601–7614, 2024.

726 Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed Chi, Sharan Narang, Aakanksha Chowdh-
 727 ery, and Denny Zhou. Self-consistency improves chain of thought reasoning in language models.
 728 *arXiv preprint arXiv:2203.11171*, 2022.

729

730 Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny
 731 Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. *NeurIPS*, 35:
 732 24824–24837, 2022.

733 Liang Wen, Yunke Cai, Fenrui Xiao, Xin He, Qi An, Zhenyu Duan, Yimin Du, Junchen Liu, Lifu
 734 Tang, Xiaowei Lv, Haosheng Zou, Yongchao Deng, Shousheng Jia, and Xiangzheng Zhang.
 735 Light-r1: Curriculum sft, DPO and RL for long COT from scratch and beyond. *arXiv preprint
 736 arXiv::2503.10460*, 2025.

737

738 Jianhao Yan, Yafu Li, Zican Hu, Zhi Wang, Ganqu Cui, Xiaoye Qu, Yu Cheng, and Yue Zhang.
 739 Learning to reason under off-policy guidance. *arXiv preprint arXiv:2504.14945*, 2025.

740 An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang
 741 Gao, Chengan Huang, Chenxu Lv, Chujie Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng Hu,
 742 Hao Ge, Haoran Wei, Huan Lin, Jialong Tang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin
 743 Yang, Jiaxi Yang, Jing Zhou, Jingren Zhou, Junyang Lin, Kai Dang, Keqin Bao, Kexin Yang,
 744 Le Yu, Lianghao Deng, Mei Li, Mingfeng Xue, Mingze Li, Pei Zhang, Peng Wang, Qin Zhu, Rui
 745 Men, Ruize Gao, Shixuan Liu, Shuang Luo, Tianhao Li, Tianyi Tang, Wenbiao Yin, Xingzhang
 746 Ren, Xinyu Wang, Xinyu Zhang, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yinger
 747 Zhang, Yu Wan, Yuqiong Liu, Zekun Wang, Zeyu Cui, Zhenru Zhang, Zhipeng Zhou, and Zihan
 748 Qiu. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*, 2025a.

749 Ling Yang, Zhaochen Yu, Bin Cui, and Mengdi Wang. Reasonflux: Hierarchical llm reasoning via
 750 scaling thought templates. *arXiv preprint arXiv:2502.06772*, 2025b.

751 Yang Yue, Zhiqi Chen, Rui Lu, Andrew Zhao, Zhaokai Wang, Yang Yue, Shiji Song, and Gao Huang.
 752 Does reinforcement learning really incentivize reasoning capacity in llms beyond the base model?
 753 *arXiv preprint arXiv:2504.13837*, 2025.

754

755 Eric Zelikman, Yuhuai Wu, Jesse Mu, and Noah D Goodman. Star: Self-taught reasoner bootstrapping
 756 reasoning with reasoning. In *NeurIPS*, volume 1126, 2024.

756 Dan Zhang, Sining Zhoubian, Ziniu Hu, Yisong Yue, Yuxiao Dong, and Jie Tang. Rest-mcts*: Llm
757 self-training via process reward guided tree search. *NeurIPS*, 37:64735–64772, 2024.
758

759 Jiajie Zhang, Nianyi Lin, Lei Hou, Ling Feng, and Juanzi Li. Adapthink: Reasoning models can
760 learn when to think. *arXiv preprint arXiv:2505.13417*, 2025a.

761 Xiaojiang Zhang, Jinghui Wang, Zifei Cheng, Wenhao Zhuang, Zheng Lin, Minglei Zhang, Shaojie
762 Wang, Yinghan Cui, Chao Wang, Junyi Peng, et al. Srpo: A cross-domain implementation of
763 large-scale reinforcement learning on llm. *arXiv preprint arXiv:2504.14286*, 2025b.
764
765
766
767
768
769
770
771
772
773
774
775
776
777
778
779
780
781
782
783
784
785
786
787
788
789
790
791
792
793
794
795
796
797
798
799
800
801
802
803
804
805
806
807
808
809

810	CONTENTS OF APPENDIX	
811		
812	A The Use of Large Language Models (LLMs)	17
813		
814	B Related Work	17
815		
816	C Limitations	18
817		
818	D Algorithm and Related Discussion	18
819		
820	D.1 The pseudocode of ReMix	18
821	D.2 Theoretical Discussion on The Historical Policy Distribution ν	18
822	D.3 Theoretical Discussion on The KL-convex Loss	19
823	D.4 The Derivation of Equation 6	19
824	D.5 Sensitivity Analysis of The Choices of Policy Reincarnation Trigger Step T	20
825		
826	E Advantage Estimation	20
827		
828	F A Brief Overview of Baseline Models	21
829		
830	F.1 1.5B Models	21
831	F.2 7B Models	21
832	F.3 Exclusion Rationale for Off-Policy Baselines	22
833		
834	G System Prompt	22
835		
836	H Key Observations from Figure 1: Efficiency–Accuracy	22
837		
838	I Training Details	23
839		
840	J Extended Evaluation of AIME’24 and AMC’23	23
841		
842	K More Training Curves	25
843		
844	L Various Analysis	26
845		
846	L.1 The Performance under Constrained Maximum Response Length	27
847	L.2 The Impact of Guide Tokens in Prompt Template	28
848		
849	M ReMix for Llama-series Base Model	29
850		
851	N Case Study	29
852		
853		
854		
855		
856		
857		
858		
859		
860		
861		
862		
863		

864 A THE USE OF LARGE LANGUAGE MODELS (LLMs)
865866 We used LLMs only for coding and writing assistance. During experimentation, we consulted the
867 LLM for code debugging. All algorithmic designs, implementations, and results were produced
868 and verified by the authors, and LLM suggestions were reviewed and tested. After completing the
869 manuscript, we used LLMS solely to polish the language (grammar and phrasing) without generating
870 new scientific content. The authors remain fully responsible for the paper’s contents.
871872 B RELATED WORK
873874 Post-training enhancement of LLM reasoning capabilities predominantly follows two paradigms (Li
875 et al., 2025b). The first, inference-time optimization, improves reasoning without updating model
876 parameters through techniques like Chain-of-Thought (CoT) prompting (Wei et al., 2022), parallel
877 reasoning and integration (Wang et al., 2022), self-reflection (Ji et al., 2023), tree-based search (Zhang
878 et al., 2024), and macro-action-guided cognitive reasoning (Liu et al., 2025a). Despite their effec-
879 tiveness, the performance of these methods is fundamentally constrained by the model’s inherent
880 capabilities. The second paradigm, parameter fine-tuning, aims to enhance these intrinsic abilities
881 into LLM. While SFT on high-quality reasoning data is a common approach, its effectiveness is often
882 limited by data availability and scalability (Zelikman et al., 2024). Consequently, RLVR has emerged
883 as a powerful alternative, learning directly from reward signals to unlock superior performance, as
884 demonstrated by models like DeepSeek-R1 (Guo et al., 2025). Notably, this differs from preference-
885 based RL which learns from a reward model trained on human/AI feedback (Bai et al., 2022a,b; Liu
886 et al., 2024), as the RLVR here utilizes direct, verifiable reward signals. Our work is situated within
887 the RFT paradigm, especially under verifiable reward.
888889 The majority of existing RFT research has relied on on-policy RL algorithms prized for their
890 training stability, such as PPO (Schulman et al., 2017). Some recent approaches have sought to
891 improve efficiency by modifying the RL architecture (e.g., GRPO (Shao et al., 2024)) or relaxing
892 optimization constraints (Seed et al., 2025). However, these on-policy RL methods exhibit severe
893 sample inefficiency, as they require fresh samples for each iteration of gradient updates. To alleviate
894 this, recent research has begun to incorporate off-policy data in RL training. Tang et al. (2025)
895 propose AGRO for a unified algorithm to leverage any-generation data, encompassing both on- and
896 off-policy samples. However, their experimental results show that off-policy training is inferior to
897 on-policy training, underscoring the non-trivial challenge of achieving stable and effective off-policy
898 training for LLMs. Tapered Off-Policy REINFORCE (Roux et al., 2025) introduces a novel variant
899 of importance sampling to downweight negative trajectories that are not likely under the current
900 policy, while allowing positive trajectories to be upweighted. This enables the utilization of both
901 off-policy and on-policy rollout trajectories. The method is trained and evaluated on GSM8K and
902 MATH, leaving its efficacy on broader reasoning tasks unknown.
903904 Recently, concurrent to our work, Based on REINFORCE, AsymRE (Arnal et al., 2025) is proposed
905 to leverage both off-policy and on-policy data by introducing a tunable baseline. An asymmetry is
906 presented that while on-policy updates safely leverage both positive and negative signals, off-policy
907 updates benefit more from positive rewards, which to some extent echoes the idea proposed in (Roux
908 et al., 2025). AsymRE is trained and evaluated on MATH. SRPO (Zhang et al., 2025b) builds on
909 GRPO with a two-stage curriculum that first trains on math and then on code. Besides, SRPO adopts
910 historical resampling discards groups with uniform rewards to avoid zero gradients and retains hard
911 samples for later replay. Using the same base model, SRPO outperforms DeepSeek-R1-Zero-Qwen-
912 32B while using only one-tenth of the training steps. Similarly, RePO (Li et al., 2025a) also exploits
913 historical data. SRPO’s resampling chiefly targets the quality of samples, while RePO emphasizes
914 efficiency and systematically analyzes the impact of replay strategies. RePO is proposed upon GRPO
915 to replay both historical off-policy data and on-policy data together during typical GRPO training.
916 Different off-policy data replay strategies are studied, among which recency-based and reward-based
917 strategies show improved performance. The RePO models are trained with a maximum response
918 length of 1,024, thus showing limited performance on math reasoning benchmarks.
919920 By following the principle of Soft RL, SPO (Cohen et al., 2025)) is proposed to leverage both off-
921 policy and on-policy data based on Cumulative Q-Parameterization. SPO is trained and evaluated for
922 code contests and demonstrates superior performance to the standard PPO. In contrast, LUFFY (Yan
923

918 et al., 2025) uses off-policy samples from superior models (e.g., DeepSeek-R1) and employing policy
 919 shaping. However, in essence, this is more akin to learning from demonstrations rather than the
 920 canonical off-policy RL where the behavior policy is often one of the historical policies or a separate
 921 inferior policy. Moreover, the idea of off-policy guidance is orthogonal to our method.

922 While early efforts have conducted first-step explorations on realizing off-policy learning for RFT,
 923 they have primarily focused on adapting existing on-policy methods (e.g., PPO, GRPO, REINFORCE)
 924 to off-policy data from the angles of modifying importance sampling, leveraging data or trajectories
 925 asymmetrically, etc. These initial steps have not investigated the essential effects of off-policy learning
 926 on reasoning behaviors, while leaving the potential of existing off-policy RL techniques unexplored.
 927 In the broader field of RL, methods such as Rainbow (Hessel et al., 2018), TD3 (Fujimoto et al.,
 928 2018), and SAC (Haarnoja et al., 2018) have set a precedent for leveraging historical data to improve
 929 sample efficiency. Building on this, advanced research has pursued maximizing data utilization
 930 through high UTD ratios, managing the resultant estimation errors with techniques like ensemble
 931 learning, as seen in REDQ (Chen et al., 2021), DroQ (Hiraoka et al., 2021), and SPEQ (Romeo et al.,
 932 2021). Concurrently, novel approaches have emerged, including hybrid methods that seek an optimal
 933 balance between the stability of on-policy learning and the efficiency of off-policy methods (Queeney
 934 et al., 2021), as well as fully offline algorithms designed to mitigate extrapolation errors from static
 935 datasets (Ma et al., 2024b; 2023). The value of ReMix lies in its departure from simply implementing
 936 off-policy RL in the context of RFT. Instead, by drawing inspiration from rich RL literature, our
 937 research aims to conduct an in-depth investigation of different off-policy RL techniques and integrate
 938 them to improve the RFT process effectively, thereby significantly enhancing the efficiency and
 939 performance of LLM fine-tuning.

C LIMITATIONS

943 Due to resource constraints, our experiments were limited to models up to 7B. While this provides
 944 a strong proof-of-concept, performance on larger-scale models is yet to be explored. To support
 945 this future work, we provide open-source code and models for community validation. We also note
 946 that the principle of using off-policy data to improve sample efficiency is general and not inherently
 947 tied to model scale. For the utilization of off-policy data, we use fixed proportions in this work,
 948 although we believe an adaptive control on the proportion of off-policy data should be possible and
 949 favorable. Moreover, our method is orthogonal to many of the advanced RFT methods considered
 950 and not considered in our experiments, while we do not explore the combination of them. We believe
 951 that integrating off-policy learning and other advanced techniques is promising to realize new LLM
 952 models that are more efficient and powerful at the same time. We leave these potential angles for the
 953 future.

D ALGORITHM AND RELATED DISCUSSION

D.1 THE PSEUDOCODE OF REMIX

954 The pseudocode of ReMix is presented in Algorithm 1.

D.2 THEORETICAL DISCUSSION ON THE HISTORICAL POLICY DISTRIBUTION ν

955 The distribution of policy index i , where $i \in \{0, 1, \dots, N\}$, is denoted by ν . The theoretical
 956 explanation for the influence of different choices of ν should trace back to how Trust-Region Policy
 957 Optimization (TRPO) (Schulman et al., 2015) approximates the trust region.

958 For $i = 0$, i.e., the standard PPO (on-policy case), the trust region $\alpha^2 = [\max_s TV(\pi_k, \pi)(s)]^2$ or its
 959 upper bound $D_{KL}(\pi_k, \pi)$ (this is can be found in Theorem 1, Eq.8 in TRPO paper) is approximated
 960 by using the expectation regarding d^{π_k} to replace the maximum case. In turn, the approximate trust
 961 region is $\tilde{\alpha}^2 = [\mathbb{E}_{s \sim d^{\pi_k}} TV(\pi_k, \pi)(s)]^2$. In the practices of PPO and TRPO, this approximation
 962 works well in many problems. Intuitively, this is because the new policy π (post-update) should
 963 not differ a lot with the current policy π_k (prior-update), the distribution d^{π_k} works as an effective
 964 surrogate.

972 **Algorithm 1** Reincarnating Mix-Policy Proximal Policy Gradient Method (**ReMix**)

973 1: [Input]: Base model π_{base} , and on-policy proximal PG method \mathbb{A} (e.g., PPO, GRPO)

974 2: Set training batch size B , off-policy data portion p , UTD ratio m , historical policy window size N , policy

975 reincarnation step point T

976 3: Init the model $\pi_\theta = \pi_{\text{base}}$ and the historical policy set $\mathbb{H} = \emptyset$

977 4: **# Stage 1: Mix-policy Proximal PG Training**

978 5: **for** step $t = 1, 2, 3, \dots, T$ **do**

979 6: Sample a batch of questions $q \sim \mathcal{D}_0$ with size $(1 - p)B$ and generate fresh responses according to π_θ

980 and \mathbb{A}

981 7: Reuse historical responses from \mathbb{H} with size pB and form the mixed training batch

982 8: Save π_θ to \mathbb{H} with its responses and logprob data, drop the oldest policy if $|\mathbb{H}| > N$

983 9: Repeatedly update π_θ with the mixed training batch according to Mix-PPG and \mathbb{A} (the first row, Eq. 5)

984 for m times

985 10: **end for**

986 11: **# Stage Transition: Policy Reincarnation**

987 12: Reset the base reference model from π_{base} to π_T , and drop the historical policy set \mathbb{H}

988 13: **# Stage 2: Reincarnating On-policy Proximal PG Training**

989 14: **for** step $t = T + 1, T + 2, T + 3, \dots$ **do**

990 15: Sample a batch of questions $q \sim \mathcal{D}_0$ with size B and generate responses according to π_θ and \mathbb{A}

991 16: Construct a training batch with the fresh responses, and update π_θ according to \mathbb{A} (the second row, Eq. 5)

992 17: **end for**

993 When it moves on to the off-policy case where $i > 0$ (Queeney et al., 2021), the approximation

994 for the trust region could be no longer effective for stable policy optimization due to the increasing

995 discrepancy between d^π and $d^{\pi_{k-i}}$ as the increase of i (i.e., for older historical policies). In our work,

996 we empirically found that using too old historical policies can introduce large off-policyness which

997 makes the training unstable.

998 Our empirical observation aligns with our discussion on the theoretical explanation of the approxima-

999 tion of trust region above. Therefore, in practice, we use $N = 2$, $p = 0.4$ and a uniform distribution

1000 for off-policy data. This is equivalent to the policy index distribution $\nu = [0.6, 0.2, 0.2]$. We believe

1001 that one important future direction is to study how to replay off-policy data better instead of replaying

1002 in a uniform manner.

1003

1004 **D.3 THEORETICAL DISCUSSION ON THE KL-CONVEX LOSS**

1005 The KL-convex loss is theoretically grounded in Conservative Policy Iteration (CPI) (Ma et al.,

1006 2024a). The CPI paper proves that iteratively refining the reference policy guarantees monotonic

1007 improvement (Proposition 1: $V^{\bar{\pi}^*}(s) \geq V^{\bar{\pi}}(s)$) and support preservation ($\bar{\pi}^*(a|s) = 0$ wherever

1008 $\bar{\pi}(a|s) = 0$). Our convex combination $L_{\text{KLC}} = \lambda D_{\text{KL}}(\pi_\theta || \pi_{\text{base}}) + (1 - \lambda) D_{\text{KL}}(\pi_\theta || \pi_{k-1})$ directly

1009 implements this — where the π_{k-1} term acts as CPI’s dynamic reference policy to prevent OOD

1010 queries, while the π_{base} term extends the framework to preserve foundational capabilities.

1011

1012 Furthermore, CPI’s Theorem 1 identifies that multi-step actor-critic implementations suffer from

1013 high variance, and the authors explicitly recommend adding behavior regularization (their Eq. 6) to

1014 constrain policies to data support. Our decaying $\lambda(t)$ schedule dynamically balances this trade-off,

1015 shifting from conservative exploration to aggressive refinement. Thus, KLC inherits CPI’s theoretical

1016 guarantees of improvement and convergence while addressing practical stability challenges in iterative

1017 LLM fine-tuning.

1018

1019

1020 **D.4 THE DERIVATION OF EQUATION 6**

1021 We conduct the formal analysis of off-policy PPG in Section 4.4 by using a similar form of average

1022 loss as in (Fatemi et al., 2025). Here, we provide the complete derivation of Equation 6 below.

1023

1024 Starting from Equation 3, that is the definition of the loss function for Mix-PPG, we rewrite the loss

1025 function from the original per-sample expectation (i.e., $(s, a) \sim d^{\pi_{k-i}}$) form to the per-trajectory

1026 expectation form (i.e., $\tau \sim d^{\pi_{k-i}}$). This derives:

$$1028 L_{\text{Avg}}^{\text{Mix-PPG}}(\theta) = -\mathbb{E}_{i \sim \nu} \left[\mathbb{E}_{\tau \sim d^{\pi_{k-i}}} \frac{1}{H_{\tau}} \sum_{h=0}^{H_{\tau}} \min \left(r_{\theta}^{k-i}(s_h, a_h) A^{\pi_k}(s_h, a_h), \right. \right. \\ 1029 \left. \left. \text{clip} \left(r_{\theta}^{k-i}(s_h, a_h), \frac{\pi_k(a_h | s_h)}{\pi_{k-i}(a_h | s_h)} - \epsilon, \frac{\pi_k(a_h | s_h)}{\pi_{k-i}(a_h | s_h)} + \epsilon \right) A^{\pi_k}(s_h, a_h) \right) \right], \quad (7) \\ 1030 \\ 1031$$

1032 where H_{τ} is response length for trajectory τ . By dropping the clipping range (i.e., ignoring the out-
1033 of-clipping range cases which have no gradient) and simplifying the expression by omitting subscript
1034 notations, expectation notations, etc., we only keep the proportional relationship for the analysis.
1035 This then leads to: $L_{\text{Avg}}^{\text{Mix-PPG}} = -\frac{1}{H} \sum_{h=0}^H r_{\theta}^{k-i} A^{\pi_k}$, which is the exact form used in Equation 6.
1036

1037 D.5 SENSITIVITY ANALYSIS OF THE CHOICES OF POLICY REINCARNATION TRIGGER STEP T

1039 In our pipeline, T is the hand-over step that switches from the off-policy Mix-PPG phase to the
1040 on-policy phase with reference model changed too. We use it to harvest early data-efficiency from
1041 replay historical data and then let on-policy optimization continue improving stably. To address
1042 the reviewer’s concern about how the trigger step T was chosen and how sensitive performance is
1043 to this choice, we provide a controlled sweep $T \in \{25, 50, 100, 200\}$ with 500-step runs, whose
1044 experimental settings are all the same as those of the main text. The results are shown in the table
1045 below. For ($T = 25$), we report two checkpoints per run, the first time the macro average reaches
1046 ≥ 49 (“early lift”) and the final plateau.

1047 The pattern is consistent: a very small T (25) slows down performance improvement; a very large T
1048 (200) lifts early but later stalls due to amplified off-policiness; a moderate T (50–100) yields both a
1049 strong early lift and the best final averages. We also find this choice to broadly work well across our
1050 experimental cases, achieving a favorable efficiency-performance trade-off.

1051 **Table 8: Pass@1 accuracy (%) of 1.5 B model under different policy reincarnation trigger step
1052 T .**

Model	AIME’24	AMC’23	MATH500	Minerva	Olympiad	Avg.
$T = 25$ (at 350 steps)	33.33	65.06	81.20	27.57	39.26	49.29
$T = 50$ (at 75 steps)	43.33	63.86	79.60	26.84	39.41	50.61
$T = 50$ (at 500 steps)	40.00	63.86	83.00	26.84	43.70	51.52
$T = 100$ (at 225 steps)	30.00	69.88	81.80	24.63	41.78	49.62
$T = 100$ (at 325 steps)	36.67	69.88	82.00	30.15	41.78	52.09
$T = 200$ (at 100 steps)	30.00	66.27	81.00	29.41	43.26	49.99
$T = 200$ (at 500 steps)	16.67	49.40	77.00	16.67	38.67	40.76

1061 E ADVANTAGE ESTIMATION

1063 To enable stable off-policy training, we adopt a V-trace (Espeholt et al., 2018) formulation for
1064 generalized advantage estimation (GAE) (Schulman et al., 2016), which incorporates truncated
1065 importance sampling ratios to correct for policy mismatch. We first compute the temporal-difference
1066 error(TD-error) at each time step t as
1067

$$1068 \delta_t^V = r(s_t, a_t) + \gamma V(s_{t+1}) - V(s_t), \quad (8)$$

1069 and define the truncated importance sampling weight $c_t = \min \left(\bar{c}, \frac{\pi_k(a|s)}{\pi_{k-i}(a|s)} \right)$, where \bar{c} is a clipping
1070 threshold to limit the variance of the correction, we use $\bar{c} = 1$ in our implement.
1071

1072 The advantage at step t is estimated recursively using the V-trace correction as
1073

$$1074 A_t = \delta_t^V + \gamma \lambda c_t A_{t+1}, \quad (9)$$

1075 and the return-to-go is computed by combining the advantage estimate with the baseline value:
1076

$$1077 \text{RTG}_t = A_t \cdot c_t + V(s_t). \quad (10)$$

1078 This V-trace corrected GAE formulation ensures that the estimated advantages remain stable and
1079 consistent under significant off-policy drift, which is critical in our training regime involving long-
horizon trajectories and evolving policies.

1080 F A BRIEF OVERVIEW OF BASELINE MODELS
10811082 F.1 1.5B MODELS
1083

- 1084 • **Open-RS Series** (Dang & Ngo, 2025): The Open-RS series employs the GRPO algorithm
1085 to train language models, using datasets constructed by filtering and combining existing
1086 corpora. Specifically, **Open-RS1** utilizes dataset with 18,615 samples with accuracy and
1087 format rewards, **Open-RS2** incorporates dataset with 7,000 samples and shorter maximum
1088 response length while retaining the same reward functions. Compared to **Open-RS2**, **Open-
1089 RS3** replaces the accuracy reward with a cosine reward and adds an English-only instruction
1090 to the system prompt.
- 1091 • **DeepScaleR** (Luo et al., 2025): DeepScaleR is obtained via a two-phase training process
1092 with the GRPO algorithm: starting with 8k context for efficient reasoning, then scaling up to
1093 16k and 24k contexts to address more challenging problems.
- 1094 • **II-Thought** (Intelligent-Internet, 2025): Based on a systematic analysis of existing public
1095 datasets, the authors constructed a large-scale, high-quality dataset comprising over 300,000
1096 reasoning problems across multiple domains. Each sample was rigorously filtered and
1097 deduplicated. Subsequently, the models were trained on this curated dataset, using the
1098 GRPO algorithm.
- 1099 • **FastCuRL Series** (Song et al., 2025): The FastCuRL Series adopts a multi-stage training
1100 process where both context length and data complexity (defined by input prompt length) are
1101 progressively increased. Training starts with short-context and low-complexity data, then
1102 moves to longer contexts with medium and high-complexity datasets.
- 1103 • **L1 Series** (Aggarwal & Welleck, 2025): The L1 Series trains models using Length-
1104 Controlled Policy Optimization (LCPO), a method that encourages correct answers while
1105 matching a target output length specified in the prompt (measured by input prompt length).
1106 **L1-Exact** enforces exact-length generation by penalizing deviation from the target length,
1107 while **L1-Max** applies a soft maximum-length constraint, allowing shorter outputs when
1108 appropriate but discouraging overruns.
- 1109 • **AdaptThink** (Zhang et al., 2025a): AdaptThink is an RFT method that trains reasoning
1110 models to choose between two modes — Thinking and NoThinking — based on problem
1111 difficulty. It uses a constrained optimization objective to encourage NoThinking while
1112 maintaining performance, and an importance sampling strategy to balance both modes
1113 during training.

1114 F.2 7B MODELS
1115

- 1116 • **ReasonFlux-F1** (Yang et al., 2025b): ReasonFlux-F1 is an SFT model obtained by fine-
1117 tuning an R1-Distill model³ based on template-augmented reasoning trajectories collected
1118 by ReasonFlux-v1. These trajectories are first enhanced with structured templates, then
1119 transformed into a long chain-of-thought format.
- 1120 • **Light-R1** (Wen et al., 2025): Light-R1 is a multi-stage post-training framework. It be-
1121 gins with curriculum-based supervised fine-tuning (SFT) using progressively harder data,
1122 followed by Direct Preference Optimization (DPO) and an RFT process with GRPO on a
1123 filtered dataset. Light-R1-7B-DS is trained only in the second SFT stage of the frame-
1124 work. Thus, the Light-R1 7B baseline model used in our experiments is an SFT model rather
1125 than an RFT model.
- 1126 • **Skywork-OR1-Preview** (He et al., 2025): Skywork-OR1-Preview is trained on a curated
1127 dataset of math and coding problems, selected through model-aware difficulty estimation
1128 and quality filtering. The training process modifies GRPO by incorporating both offline
1129 and online difficulty-based filtering, rejection sampling, and a multi-stage curriculum with
1130 adaptive entropy control.
- 1131 • **Polaris** (An et al., 2025): Polaris adopts a multi-stage RL training approach with careful data
1132 difficulty control, using a data distribution with a slight bias toward challenging problems and

1133 ³https://github.com/Gen-Verse/ReasonFlux/blob/main/ReasonFlux_F1/README.md

1134 dynamically adjusting question difficulty during training. It initializes sampling temperature
 1135 based on rollout diversity and gradually increases it during training. It employs length
 1136 extrapolation techniques, enabling longer CoT generation at inference while keeping training
 1137 rollouts short.

1138 • **AdaptThink** (Zhang et al., 2025a): The methodology for the AdaptThink 7B model is
 1139 identical to that of the AdaptThink 1.5B model, as previously described.
 1140 • **AceReason-Nemotron** (Chen et al., 2025): AceReason-Nemotron adopts the GRPO al-
 1141 gorithm without KL divergence and avoids entropy collapse through controlled updates.
 1142 The model is first trained on math-only prompts, then on code-only prompts, following a
 1143 curriculum with progressively increasing response lengths.

1144

1145 F.3 EXCLUSION RATIONALE FOR OFF-POLICY BASELINES

1146

1147 For existing off-policy RFT methods, we do not include RePO (Li et al., 2025a) because their models
 1148 are trained under a maximum response length of 1,024 tokens, thus showing limited performance
 1149 on math reasoning tasks. We do not include LUFFY (Yan et al., 2025) since the usage of off-
 1150 policy guidance from a superior model (e.g., DeepSeek-R1) is orthogonal to ReMix, which is also
 1151 viewed as a different setting where extrinsic guidance or demonstrations are accessible. We exclude
 1152 SRPO (Zhang et al., 2025b) since its publicly released model is not at the same scale as ours. In
 1153 addition, we did not find public checkpoints for SPO (Cohen et al., 2025) (which is also trained
 1154 for code contests), AGRO (Tang et al., 2025), AsymRE (Arnal et al., 2025) and Tapered Off-policy
 1155 REINFORCE (Roux et al., 2025), thus, we do not include them in our experiments. Please refer to
 1156 Section B for detailed discussions on related off-policy RFT methods.

1157

1158 G SYSTEM PROMPT

1159

1160 Following the standard DeepScaler data processing approach, each prompt in the training set was
 1161 prefixed with "<|User|>" and suffixed with the instruction "Let's think step by step and output
 1162 the final answer within \boxed{.} <|Assistant|><think>". This structure encourages the model
 1163 to engage in step-by-step reasoning and produce final answers encapsulated within LaTeX boxed
 1164 expressions. One example of the DeepScaler prompt format is shown below. The blue text
 1165 indicates the fixed template used during inference, while the black text represents the instance-specific question
 1166 inserted into the prompt.

1167 System Prompt (Standard)

1168 <|begin_of_sentence|><|User|> Xenia and Sergey play the following game. Xenia thinks
 1169 of a positive integer N not exceeding 5000. Then she fixes 20 distinct positive integers
 1170 a_1, a_2, \dots, a_{20} such that, for each $k = 1, 2, \dots, 20$, the numbers N and a_k are congruent
 1171 modulo k . By a move, Sergey tells Xenia a set S of positive integers not exceeding 20,
 1172 and she tells him back the set $\{a_k : k \in S\}$ without spelling out which number corresponds
 1173 to which index. How many moves does Sergey need to determine for sure the number
 1174 Xenia thought of? Let's think step by step and output the final answer within \boxed{.}.

1175 <|Assistant|><think>

1176

1177

1178 H KEY OBSERVATIONS FROM FIGURE 1: EFFICIENCY–ACCURACY

1179

1180 For clarity, we summarize the major observations in Figure 1 below:

1181

- 1182 • **(1.5B) ReMix-PPO v.s., DeepScaleR:** DeepScaleR, the strongest 1.5B competitor, requires
 1183 around **2.519M** rollouts to reach its final score (i.e., **52.14**), whereas ReMix-PPO (350 Steps)
 1184 achieves a comparable score (i.e., **52.10**) with **0.079M** rollouts — over a **30x reduction** in
 1185 rollout data volume.
- 1186 • **(1.5B) ReMix-PPO v.s., PPO:** We trace the performance of ReMix-PPO at 100, 200, and
 1187 350 training steps (denoted by the **yellow curve** in Figure 1), corresponding to rollout data
 1188 volumes of roughly 0.020M, 0.041M, and 0.079M, respectively. Even after generating just

1188 **0.020M** rollout samples, ReMix-PPO achieves a score of **50.61**, which has already surpasses
 1189 most baselines. Compared to PPO (900 Steps), which achieves an average score of **50.61**
 1190 with **0.230M** rollouts, our model shows over **a 10x reduction** in rollout data volume.
 1191
 1192 • **(1.5B) ReMix-GRPO v.s., GRPO:** We also trace the performance of ReMix-GRPO at 50,
 1193 100, and 200 training steps (denoted by the **cyan curve** in Figure 1), corresponding to rollout
 1194 data volumes of roughly 0.061M, 0.163M, and 0.368M, respectively. After generating
 1195 **0.061M** rollout samples, our model achieves the score **45.47** that exceeds the score **45.20** of
 1196 standard GRPO trained for 100 steps with **0.205M** rollout samples. Compared to GRPO
 1197 (200 Steps), which achieves an average score of **48.64** with **0.410M** rollouts, ReMix-GRPO
 1198 achieves a much higher score of **50.53** within 200 training steps, i.e., **0.368M** rollouts,
 1199 showing a superior final performance with less computational cost.
 1200
 1201 • **(7B) ReMix-PPO v.s., AceReason-Nemotron:** AceReason-Nemotron, the strongest 7B
 1202 baseline method in our comparison, requires over **3.584M** rollouts to reach its final score
 1203 (i.e., **63.24**⁴), whereas ReMix-PPO (50, 75 Steps) achieves a slightly higher accuracy (i.e.,
 1204 **63.27**, **64.39**) with **0.007M**, **0.011M** rollouts — over **a 450x reduction** in rollout data
 1205 volume.
 1206
 1207 • **(7B) ReMix-PPO v.s., AdaptThink:** AdaptThink, the second strongest 7B baseline method,
 1208 requires around **0.307M** rollouts to reach its final score (i.e., **58.77**), whereas ReMix-PPO
 1209 (25 Steps) achieves a comparable accuracy (i.e., **58.49**) with **0.003M** rollouts — over **an**
 1210 **80x reduction** in rollout data volume.
 1211
 1212 Besides, when `do_sample` is set to `true`, Open-RS series models (i.e., -RS1, -RS2, -RS3) show
 1213 better scores 40.62, 40.08, 39.31 respectively, and `TI-Thought` can achieve a score 51.474. For
 1214 other models, we found similar scores in our experiments, which do not change the conclusions.

I TRAINING DETAILS

Hyperparameters The major hyperparameter choices are shown in Table 9.

Compute Resource The 1.5B model was trained for 50 hours on 2 NVIDIA A800-SXM4-80GB GPUs, while the 7B model required 75 hours on 8 such GPUs. The evaluation of each model was also conducted using the same number of GPUs as in their respective training setups.

Comparison of Training Detail on Computational Cost for 1.5B Models The corresponding detailed factors associated with computational cost for training the 1.5B models in the comparison above are shown in Table 10. Compared to most baselines, our method uses nearly half the number of training steps (500 v.s. ≥ 860) while delivering superior performance. Furthermore, our entire training run is executed on a single node with just two A800 GPUs over 52 hours, amounting to 104 A800 GPU hours. This finding shows that state-of-the-art gains can be achieved with markedly reduced compute requirements.

Comparison of Training Detail on Computational Cost for 7B Models Table 11 shows the training details of 7B models. However, we failed to find complete training details for all the 7B models, so we did not plot the efficiency-performance trade-off for the 7B models due to missing information.

J EXTENDED EVALUATION OF AIME’24 AND AMC’23

For two small and high-variance benchmarks (i.e. AIME’24 and AMC’23), we report Pass@1 as **the average over 32 independent runs (Avg@32)**. Here we use stochastic decoding (temperature 0.1) for evaluation. We include these results in Table 12 and Table 13 to provide a robust estimate on small datasets. ReMix-PPO maintains the similar performance observed in Table 1 and Table 2 under greedy decoding, showing consistent gains under the Avg@32 protocol.

⁴The score of AceReason-Nemotron is obtained by evaluating the official checkpoint, and the rollout data volume is estimated according to the text and Figure 3 in (Chen et al., 2025).

1242

1243

Table 9: **Hyperparameter setups for PPO, GRPO and ReMix trainer.**

1244

1245

1246

1247

1248

1249

1250

1251

1252

1253

1254

1255

1256

1257

1258

1259

1260

1261

1262

1263

1264

1265

1266

1267

1268

1269

1270

1271

Parameter	Value
<i>Training Configuration</i>	
temperature	1.0
top-p	1.0
top-k	-1
critic_warmup	0
learning_rate	1e-6
clip_ratio	0.2
lam	1
tau	0.95
entropy_coeff	0.001
clipping_gradient	true
do_sample	true
test_freq	25
<i>Training Configuration for ReMix-GRPO and GRPO</i>	
kl_loss_coef	0.001
kl_loss_type	low_var_kl
n (gen per prompt)	8

Table 10: **RFT training details associated with computational cost for 1.5B models.** All the models are trained upon DeepSeek-R1-distill-Qwen2.5 base model, except for L1 series models, which are fine-tuned on top of DeepScaleR (denoted by superscript *). Accordingly, their total training cost should be considered as the sum of DeepScaleR’s cost and the resources reported in this table. *Italicized entries* indicate values not directly reported in the original papers, but instead retrieved from associated official training scripts. The underlined values denote the fresh on-policy rollout in addition to off-policy data reuse in ReMix.

Model	Traing Steps	Rollout Batch Size	Gen per Prompt	Max Responses Length	Number of GPUs
DeepScaleR	1750 steps	<i>128,128,128</i>	<i>8,16,16</i>	8k,16k,24k	8,32,32
FASTCuRL-preview	860 steps	128,64,64,64	8,8,8,16	8k,16k,24k,16k	8
FASTCuRL-v3	2620 steps	<i>128,64,64,64,64</i>	<i>8,8,8,16,16</i>	8k,16k,24k,16k,16k	8
II-Thought	-	<i>1024</i>	<i>5</i>	32k	8
adapt think	314 steps	128	16	16k	8
Open-RS1	100 steps	96	6	4k	4
Open-RS2	50 steps	96	6	4k	4
Open-RS3	50 steps	96	6	4k	4
L1-Exact*	700 steps	128	16	4k	8
L1-Max*	120 steps	128	16	4k	8
ReMix-PPO	500 steps	<u>152,256</u>	1	8k	2
ReMix-GRPO	200 steps	<u>152,256</u>	8	8k	2

Table 11: **RFT training details associated with computational cost for 7B models.** All methods are trained upon DeepSeek-R1-Distill-Qwen-7B base model. *Italicized entries* indicate values not directly reported in the original papers, but instead retrieved from associated official training scripts. The underlined values denote the fresh on-policy rollout in addition to off-policy data reuse in ReMix. Note that ReasonFlux-F1 and Light-R1 (7B) are SFT models as detailed in Appendix F.2, hence we do not include them in this table.

Model	Traing Steps	Rollout Batch Size	Gen per Prompt	Maximum Responses Length	Number of GPUs
Skywork-OR1-Preview	>2000 steps	256	<i>16</i>	8k,16k,32k	8
AceReason-Nemotron	>2000 steps	128	8,16,16,16	8k,16k,24k,32k	128
AdaptThink	150 steps	128	16	16k	8
Polaris	>1400 steps	-	-	16k,24k,32k	-
ReMix-PPO	500 steps	<u>152,256</u>	1	8k	8

1296 **Table 12: Pass@1 accuracy (%) on AIME’24 and AMC’23 (Avg@32) of 1.5B models.** ReMix-
 1297 PPO shows performance consistent with the greedy-decoding results (Table 1). **Bolded** and underlined
 1298 values denote the highest and the second-highest scores in each dataset (i.e., column).

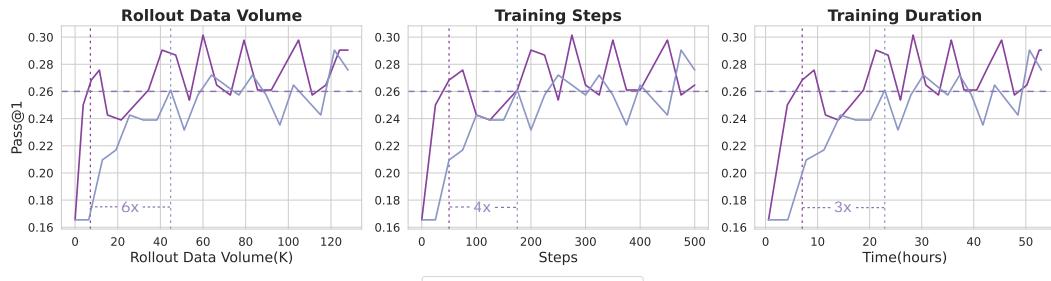
Model	AIME’24	AMC’23
R1-Distill-Qwen-1.5B (Base)	17.92	44.58
Open-RS1	17.60	45.44
Open-RS2	17.62	44.92
Open-RS3	18.65	43.98
AdaptThink	19.06	56.81
II-Thought	28.65	59.72
FASTCuRL-preview	25.94	54.27
FASTCuRL-V3	33.44	63.63
L1-Exact*	25.1	<u>66.57</u>
L1-Max*	23.13	66.79
DeepScaleR	<u>31.96</u>	63.58
ReMix-GRPO (75 Steps)	25.10	60.17
ReMix-PPO (350 Steps)	29.08	64.04

1314 **Table 13: Pass@1 accuracy (%) on AIME’24 and AMC’23 (Avg@32) of 7B models.** ReMix-PPO
 1315 reach SOTA-level on both benchmarks.

Model	AIME’24	AMC’23
R1-Distill-Qwen-7B (Base)	37.53	66.55
ReasonFlux-F1	20.19	53.07
Light-R1	40.00	66.73
Skywork-OR1-Preview	36.31	61.60
Polaris	39.71	67.40
AdaptThink	47.62	74.20
AceReason-Nemotron	<u>50.00</u>	77.48
ReMix-GRPO (200 Steps)	49.48	82.12
ReMix-PPO (75 Steps)	50.31	77.78

K MORE TRAINING CURVES

1329 **Training Curves for Efficiency Comparison** In addition to the efficiency comparison between
 1330 ReMix-PPO and PPO for Olympiad in Figure 3, the remaining curves for the other four math
 1331 reasoning benchmarks are presented in Figure 5, 6, 7, 8.



1343 **Figure 5: Training efficiency comparison for ReMix-PPO and PPO (1.5B) on Minerva.** ReMix
 1344 achieves a score above 26%, around 3x to 6x faster than PPO.

1345 **Training Curves under Varying Proportions of Off-policy Data** We vary the off-policy proportion
 1346 $p \in \{0.1, 0.2, 0.3, 0.4, 0.5\}$ (the UTD ratio is set to 1 here for isolation) to isolate the effect of
 1347 historical data reuse. In addition to the Pass@1 accuracy, we use three more metrics: the importance
 1348 sampling ratio r_{θ}^{k-i} , ratio $\frac{\pi_k}{\pi_{k-i}}$ that quantifies the distributional shift between current and historical
 1349 policies, and the response length that reflects the reasoning behavior of the model.

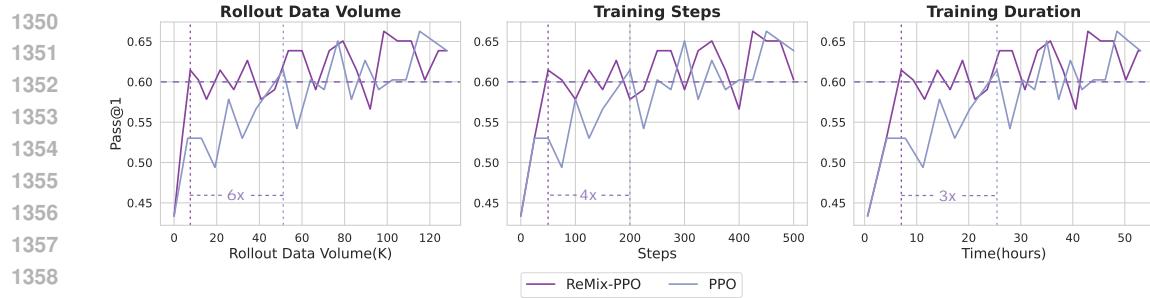


Figure 6: **Training efficiency comparison for ReMix-PPO and PPO (1.5B) on AMC'23.** ReMix achieves a score above 60%, around **3x to 6x** faster than PPO.

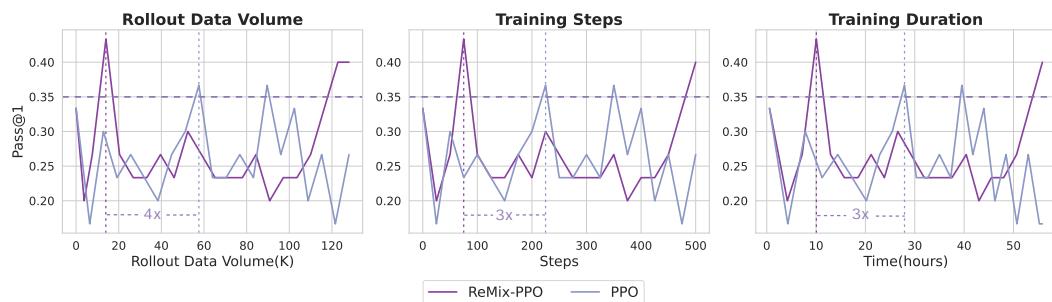


Figure 7: **Training efficiency comparison for ReMix-PPO and PPO (1.5B) on AIME'24.** ReMix achieves a score above 35%, around **1.2x to 1.6x** faster than PPO.

For importance sampling ratio, we observe that larger p yields slightly wilder importance-sampling swings. Notably, the consistently slight decrease of the importance sampling in Figure 9 can also be explained by the shortening of response length, as the whipping effect (detailed in Subsection 4.4) gradually diminishes.

Training Curves for Policy Loss Figure 10 shows that during the training process, the policy loss predominantly remains positive, which means a larger importance ratio will lead to a larger policy loss.

L VARIOUS ANALYSIS

In this section, we further analysis the effects of ReMix, including the performance under response length constraint, the impact of prompt format.



Figure 8: **Training efficiency comparison for ReMix-PPO and PPO (1.5B) on MATH500.** ReMix achieves a score above 80%, around **3x to 5x** faster than PPO.

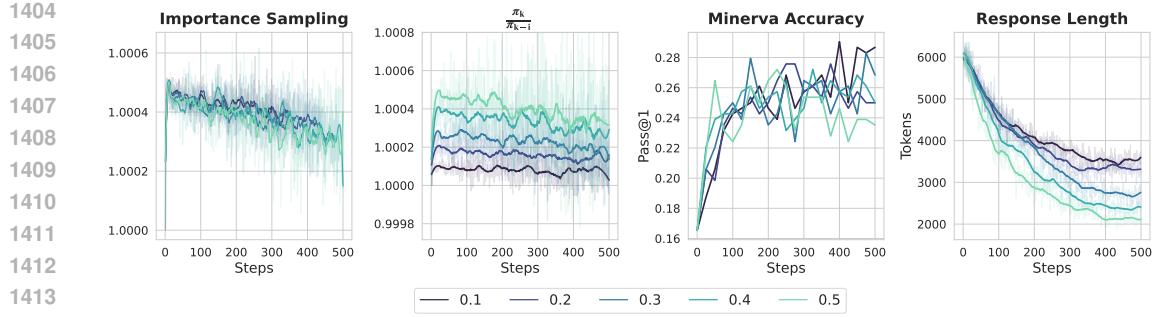


Figure 9: **Training dynamics regarding importance sampling ratio, accuracy, and response length under varying proportions of off-policy data p for Mix-PPG.** Leveraging more off-policy data leads to a larger policy distribution shift, a faster early boost in accuracy yet worse later-stage performance, and a shorter response length.

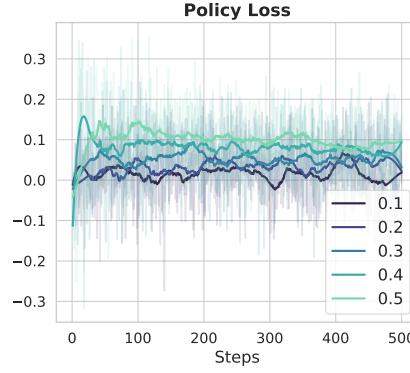


Figure 10: **Policy loss under varying proportions of off-policy data p for Mix-PPG.** Leveraging more off-policy data leads to larger policy loss.

L.1 THE PERFORMANCE UNDER CONSTRAINED MAXIMUM RESPONSE LENGTH

Since ReMix shows a feature in generating more concise responses as discovered above, we conduct an additional experiment to evaluate the performance of our model when the maximum response length is constrained. Different from the default evaluation setting of 8,192 maximum response length, we halve the maximum response length to 4,196 tokens for ReMix during evaluation. For comparison, we evaluate ReMix-PPO (1.5B) with the base model, DeepScaleR, and PPO under the halved maximum response length. The results are summarized in Table 14.

Table 14: **Performance evaluation of 1.5B models with 4k maximum response length.** The arrow \downarrow denotes the accuracy degradation compared to the results with 8k maximum response length (referring to the results in Table 1). All the models are negatively influenced by the halved maximum response length. Compared with DeepScaleR, **ReMix-PPO exhibits the smallest decrease in model performance and performs the best**, thanks to its concise and shorter reasoning behaviors.

Model	AIME'24	AMC'23	MATH500	Minerva	Olympiad	Avg.
Maximum response length: 4096 tokens						
R1-Distill-Qwen-1.5B (Base Model)	20.00	37.35	60.40	13.24	22.37	30.67 _{16.91}
DeepScaleR	10.00	49.4	75.00	21.32	34.22	37.99 _{14.15}
PPO (500 Steps)	20.00	48.19	77.60	25.00	38.96	41.95 _{17.61}
ReMix-PPO (350 Steps)	23.33	59.04	79.00	27.57	39.11	45.61_{16.49}

All the models are negatively influenced by the halved maximum response length, which matches the intuition. Notably, DeepScaleR, the best 1.5B baseline model used in our work, suffers a significant

1458 performance drop when the maximum response length is limited to 4,192 tokens. In contrast, ReMix-
 1459 PPO exhibits the smallest decrease in model performance and performs the best in this constrained
 1460 setting. This finding underscores the resilience of the preference for a concise and shorter reasoning
 1461 process learned via off-policy training of ReMix in handling the constraints on response length.
 1462

1463 **Takeaway 4. ReMix favors concise reasoning and is resilient to constraints on response
 1464 length.**

1465 The constraint on maximum response length greatly degrades the test-time reasoning per-
 1466 formance of LLM models, while ReMix suffers less thanks to its preference for a concise
 1467 reasoning process.
 1468

1470 **L.2 THE IMPACT OF GUIDE TOKENS IN PROMPT TEMPLATE**

1471 In addition, we investigate the critical role of the prompt template for response generation used
 1472 during training and evaluation. To establish a comparison, we make use of a prompt template
 1473 *without guide tokens* (as shown below). Recall the standard prompt template we presented in
 1474 Section 4.1, the difference is that the prompt template without guide tokens does not contain the
 1475 guide tokens that appear as the prefix (i.e., <begin_of_sentence><|User|>) and the suffix
 1476 (i.e., <|Assistant|><think>).
 1477

1478 **System Prompt (Without Guide Tokens)**

1479 `<|begin_of_sentence|><|User|>Xenia and Sergey play the following game. Xenia thinks
 1480 of a positive integer N not exceeding 5000. Then she fixes 20 distinct positive integers
 1481 a_1, a_2, \dots, a_{20} such that, for each $k = 1, 2, \dots, 20$, the numbers N and a_k are congruent
 1482 modulo k . By a move, Sergey tells Xenia a set S of positive integers not exceeding 20,
 1483 and she tells him back the set $\{a_k : k \in S\}$ without spelling out which number corresponds
 1484 to which index. How many moves does Sergey need to determine for sure the number
 1485 Xenia thought of? Let's think step by step and output the final answer within \boxed{ }.`
 1486 `<|Assistant|><think>`

1487 To investigate the impact of different prompt templates, we define a response as format-correct if
 1488 it includes content enclosed within paired <think> </think>. Parallel to the 1.5B base model
 1489 and ReMix-PPO, we consider a variant of ReMix-PPO that is trained without guide tokens, denoted
 1490 as ReMix-PPO w/o Guide Tokens. We evaluate the performance of the candidate models in terms
 1491 of Pass@1 accuracy⁵ and format correctness on MATH500, when using the standard template (i.e.,
 1492 with guide tokens) and the modified template without guide tokens. The purpose of this experiment
 1493 is to answer two questions: (1) whether the models trained with guide tokens (i.e., the base model,
 1494 ReMix-PPO) can also perform well when the guide tokens are not prompted during evaluation; (2)
 1495 whether the model trained without guide tokens can also obey the format and output the solution. The
 1496 results are summarized in Table 15.
 1497

1498 **Table 15: Performance evaluation of 1.5B models with and without guide tokens on MATH500.**
 1499 Both the base model and ReMix-PPO show a 0 format correctness when the guide tokens are missing
 1500 during evaluation, while ReMix-PPO exhibits a smaller drop in the accuracy. For the variant of
 1501 ReMix trained without the guide tokens, it **performs well under both the two template settings.** ↑
 1502 means higher is better and ↓ means lower is better.

Model	Eval w/ Standard Temp. (↑)		Eval w/o Guide Tokens (↑)		Relative Decrease (↓)	
	Pass@1	Format Cor.	Pass@1	Format Cor.	Pass@1	Format Cor.
R1-Distill-Qwen-1.5B (Base Model)	67.40	70.00	52.00	0	15.40	70.00
ReMix-PPO (350 Steps)	82.00	93.60	71.00	0	11.00	93.60
ReMix-PPO w/o Guide Tokens (500 Steps)	82.00	92.20	77.60	91.20	4.40	1.00

1511 ⁵Note that the correct answer with a wrong format is still counted as correct for Pass@1 accuracy here.

1512 The results show that the base model yields a format correctness of 0 when evaluated without the
 1513 guide tokens, accompanied by a decrease of 15.40 points in Pass@1 accuracy. Similarly, ReMix-PPO
 1514 also exhibits a 0 format correctness yet a smaller decrease of 11.00 in the accuracy. This indicates
 1515 that the presence of the <think> token in the prompt helps the model to autonomously generate a
 1516 closing </think> tag, maintaining format consistency. Thus, it delivers a negative answer to the
 1517 first question above, while ReMix shows a better robustness to the absence of the guide tokens.

1518 In contrast, the variant trained without the guide tokens also performs well when using the standard
 1519 template, and achieves an increase from 77.60 to 82.00, reaching the same performance as ReMix
 1520 that is trained with the guide tokens explicitly. It also maintains consistently high format correctness.
 1521 This shows a good robustness to prompt change. We found similar results for the other four math
 1522 reasoning tasks as well.

1523 We hypothesize that removing the guide tokens during the training of ReMix allows the model to
 1524 explore a broader distribution, rather than overfitting to the explicit guide tokens in the standard
 1525 template. Such flexibility encourages the model to internalize reasoning behavior in a robust and
 1526 general manner, instead of relying on external structural cues too much. As a result, it becomes
 1527 more robust to prompt variation at inference time. The smaller relative degradation observed in both
 1528 accuracy and format correctness supports this view.

1529 **Takeaway 5. ReMix is more robust to the variation of prompt template.**

1530
 1531 Removing explicit guide tokens in the standard template significantly cripples the performance
 1532 of the base model, while ReMix exhibits better robustness and compatibility to the absence
 1533 of the guide tokens during both training and evaluation.
 1534

1535
 1536 **M REMIX FOR LLAMA-SERIES BASE MODEL**
 1537

1538
 1539 **Table 16: Pass@1 accuracy (%) of R1-Distill-Llama-8B.**

Model	AIME'24	AMC'23	MATH500	Minerva	Olympiad	Avg.	Cost
R1-Distill-Llama-8B (Base Model)	26.46	62.69	82.00	25.37	42.67	47.84	N/A
PPO (50 steps)	34.58	70.63	85.40	26.10	45.18	52.38	0.013 M
ReMix-PPO (25 steps)	39.48	88.00	87.04	28.68	48.74	55.99	0.004 M

1540
 1541 **N CASE STUDY**
 1542

1543 To better understand model’s reasoning behavior, we present a case study centered on a representative
 1544 example that the base model is able to solve correctly. Figure compares the responses produced by
 1545 three variants trained under distinct strategies: PPO, Mix-PPG, and Mix-PPG with an Increased UTD
 1546 ratio. Notably, the three outputs differ significantly in length, with the PPO-trained model producing
 1547 the longest response, followed by Mix-PPG, and Mix-PPG with an Increased UTD ratio yielding the
 1548 shortest.

1549 With a sufficiently long response window, the model engages in explicit self-reflection and follows
 1550 a structured step-by-step reasoning process to arrive at the correct answer. In contrast, the UTD-2
 1551 model, exhibits minimal or no reflective behavior and tends to bypass intermediate reasoning steps,
 1552 leading to a more direct but less interpretable answer. These observations suggest that adequate
 1553 response length plays a critical role in enabling reflective, multi-step reasoning.

1554
 1555
 1556
 1557
 1558
 1559
 1560
 1561
 1562
 1563
 1564
 1565

1566
1567**Question**

1568

Prompt

1569

Five points A , B , C , D , O , H , W , X , Y , Z lie on a flat field. A is directly north of O , B is directly west of O , C is directly south of O , and D is directly east of O . The distance between C and D is 140 m. A hot-air balloon is positioned in the air at H directly above O . The balloon is held in place by four ropes HA , HB , HC , and HD . Rope HC has length 150 m and rope HD has length 130 m. [asy]
size(250);
pair A, B, C, D, O, H, W, X, Y, Z;
O=(0,0);
A=(1,1);
D=(1.5,-.3);
B=(-1.5,.3);
C=(-1,-1);
H=(0,2.5);
W=(5/3)*(A+D);
X=(5/3)*(A+B);
Y=(-1)*(W);
Z=(-1)*(X);
draw(W--X--Y--Z--W);
draw(A--C);
draw(B--D);
draw(O--H, linewidth(1));
draw(A--H, dashed);
draw(B--H, dashed);
draw(C--H, dashed);
draw(D--H, dashed);
dot(A);
dot(B);
dot(C);
dot(D);
dot(O);
dot(H);
label("A", A, NE);
label("B", B, SW);
label("C", C, SE);
label("D", D, NE);
label("O", O, SE);
label("H", H, NW);
[/asy]

To reduce the total length of rope used, rope HC and rope HD are to be replaced by a single rope HP where P is a point on the straight line between C and D . (The balloon remains at the same position H above O as described above.) Determine the greatest length of rope that can be saved.

1597

Correct Answer

1598

160

1600

1601

1602

1603

1604

1605

1606

1607

1608

1609

1610

1611

1612

1613

1614

1615

1616

1617

1618

1619

1620

Case 1: PPO (1/5)

1621

1622

Stage 1 - Question Rehearsal

1623

Okay, so I have this problem with five points on a flat field: A, B, C, D, and O. A is directly north of O, B is directly west of O, C is directly south of O, and D is directly east of O. The distance between C and D is 140 meters. There's a hot-air balloon at point H directly above O. The balloon is held by four ropes: HA, HB, HC, and HD.

1624

Now, the problem says that rope HC is 150 meters long and rope HD is 130 meters long. We need to reduce the total length of rope used by replacing HC and HD with a single rope HP, where P is a point on the straight line between C and D. The balloon stays at the same position H above O. We need to find the greatest length of rope that can be saved.

1625

Stage 2 - Spatial Configuration and Equation Derivation

1626

First, I need to visualize the setup. Points A, B, C, D are all around O, which is the center. A is north, B is west, C is south, D is east. So, if I imagine a coordinate system with O at $(0,0)$, then A would be at $(0, a)$, B at $(-b, 0)$, C at $(0, -c)$, and D at $(d, 0)$. But wait, actually, since A is directly north of O, B is directly west, C is directly south, and D is directly east, the coordinates can be set as follows:

1627

Let's assume O is at $(0,0)$. Then, A is at $(0, h)$, B is at $(-k, 0)$, C is at $(0, -m)$, and D is at $(n, 0)$. But actually, the distances between C and D is given as 140 meters. Since C is directly south of O and D is directly east of O, the distance between C and D would be the hypotenuse of the triangle formed by OC and OD. So, OC is the distance from O to C, which is m , and OD is the distance from O to D, which is n . Therefore, the distance between C and D is $\sqrt{n^2 + m^2} = 140$ meters.

1628

But we don't know m and n yet. However, we know the lengths of HC and HD. Let me recall that H is directly above O, so its coordinates are $(0, h)$. The distance from H to C is 150 meters, and the distance from H to D is 130 meters.

1629

Stage 3 - Step-by-Step Computation

1630

So, let's write the distance formulas:

1631

Distance HC: $\sqrt{(0 - 0)^2 + (h - (-m))^2} = \sqrt{(h + m)^2} = |h + m| = 150$. Since h and m are positive distances, $h + m = 150$.

1632

Similarly, distance HD: $\sqrt{(0 - n)^2 + (h - 0)^2} = \sqrt{n^2 + h^2} = 130$. So, $\sqrt{n^2 + h^2} = 130$.

1633

We also know that the distance between C and D is 140 meters, which is $\sqrt{n^2 + m^2} = 140$.

1634

So, we have three equations:

1635

1. $h + m = 150$
2. $\sqrt{n^2 + h^2} = 130$
3. $\sqrt{n^2 + m^2} = 140$

1636

We can square the third equation to get $n^2 + m^2 = 140^2 = 19600$.

1637

From the first equation, $h = 150 - m$.

1638

Substitute h into the second equation:

1639

$$\sqrt{n^2 + (150 - m)^2} = 130$$

1640

Square both sides:

1641

$$n^2 + (150 - m)^2 = 130^2 = 16900$$

1642

1643

1644

1645

1646

1647

1648

1649

1650

1651

1652

1653

1654

1655

1656

1657

1658

1659

1660

1661

1674
1675 **Case 2: PPO (2/5)**
1676 We also have $n^2 + m^2 = 19600$.
1677 Let's subtract the second equation from the first:
1678 $(n^2 + (150 - m)^2) - (n^2 + m^2) = 16900 - 19600$
1679
1680 Expanding $(150 - m)^2$:
1681 $n^2 + 22500 - 300m + m^2 - n^2 - m^2 = -2700$
1682
1683 Simplify:
1684 $22500 - 300m = -2700$
1685
1686 So, $22500 + 2700 = 300m$
1687
1688 Divide both sides by 300:
1689 $m = 25200 / 300 = 84$
1690
1691 So, $m = 84$ meters.
1692
1693 Then, from the first equation, $h = 150 - m = 150 - 84 = 66$ meters.
1694
1695 Now, from the third equation, $n^2 + m^2 = 19600$
1696
1697 We know $m = 84$, so $n^2 + 84^2 = 19600$
1698
1699 Calculate 84^2 : $84*84 = 7056$
1700
1701 So, $n^2 = 19600 - 7056 = 12544$
1702
1703 Therefore, $n = \sqrt{12544} = 112$ meters.
1704
1705 So, we have:
1706

- $m = 84$ meters (distance from O to C)
- $n = 112$ meters (distance from O to D)
- $h = 66$ meters (height of H above O)

1707
1708 **## Stage 4 - Optimization Strategy**
1709
1710 Now, we need to replace ropes HC and HD with a single rope HP, where P is a point on CD.
1711 We need to find the position of P such that the total length of HP is minimized, thus
1712 maximizing the length saved.
1713
1714 So, the original total length is $HC + HD = 150 + 130 = 280$ meters.
1715
1716 **## Stage 5 - Critical Point Analysis**
1717
1718
1719
1720
1721
1722
1723
1724
1725
1726
1727

1728
1729**Case 3: PPO (3/5)**1730
1731
1732
1733**## Stage 6 - Step-by-Step Computation**

1734 Alternatively, we can model this as a reflection of H across CD, and then the minimal
 1735 distance from H to P to D would be the straight line distance from H to the reflection
 1736 of D across CD.

1737 Wait, actually, since P is on CD, we can reflect H across CD to get a point H', and then
 1738 the minimal distance from H to P to D is the same as the distance from H' to D. But since
 1739 P is on CD, the minimal HP is the distance from H to the line CD, but I might be mixing
 1740 things up.

1741 Wait, perhaps a better approach is to parametrize point P on CD and then find the minimal
 1742 distance HP.

1743 Let's set up a coordinate system.

1744 Let's place O at (0,0). Then, C is at (0, -84), D is at (112, 0). H is at (0, 66).

1745 We need to find a point P on CD such that HP is minimized.

1746 Parametrize CD: from C (0, -84) to D (112, 0). Let's let t be the parameter from 0 to 1,
 1747 where t=0 is C and t=1 is D.

1748 So, the coordinates of P can be expressed as:

1749 $x = 112t$

1750 $y = -84 + 84t$

1751 So, $P = (112t, -84 + 84t)$

1752 Now, the distance HP is the distance between H (0,66) and P (112t, -84 + 84t).

1753 So, distance squared is:

1754 $(112t - 0)^2 + (-84 + 84t - 66)^2$

1755 Simplify:

1756 $(112t)^2 + (-150 + 84t)^2$

1757 Let's compute this:

1758 $(12544t^2) + (22500 - 2*150*84t + 7056t^2)$

1759 Compute each term:

1760 $12544t^2 + 22500 - 25200t + 7056t^2$

1761 Combine like terms:

1762 $(12544 + 7056)t^2 - 25200t + 22500$

1763 Calculate 12544 + 7056:

1764 $12544 + 7056 = 19600$

1765 So, distance squared is $19600t^2 - 25200t + 22500$

1766 To find the minimum, take derivative with respect to t and set to zero.

1767

1768

1769

1770

1771

1772

1773

1774

1775

1776

1777

1778

1779

1780

1781

Case 4: PPO (4/5)

1782 Let's denote $f(t) = 19600t^2 - 25200t + 22500$
1783
1784 $f'(t) = 39200t - 25200$
1785
1786 Set $f'(t) = 0$:
1787
1788 $39200t - 25200 = 0$
1789
1790 $39200t = 25200$
1791
1792 $t = 25200 / 39200$
1793
1794 Simplify:
1795 Divide numerator and denominator by 2800:
1796
1797 $25200 / 2800 = 9$
1798
1799 $39200 / 2800 = 14$
1800 So, $t = 9/14$
1801
1802 Therefore, the point P is at:
1803
1804 $x = 112 * (9/14) = 112 * (9) / 14 = 8 * 9 = 72$
1805
1806 $y = -84 + 84 * (9/14) = -84 + (84 * 9) / 14 = -84 + (6 * 9) = -84 + 54 = -30$
1807
1808 So, P is at (72, -30)
1809
1810 Now, compute the distance HP:
1811
1812 Distance = $\sqrt{(72 - 0)^2 + (-30 - 66)^2} = \sqrt{72^2 + (-96)^2}$
1813
1814 Compute 72^2 : 5184
1815
1816 Compute $(-96)^2$: 9216
1817
1818 Sum: $5184 + 9216 = 14400$
1819
1820 So, distance HP = $\sqrt{14400} = 120$ meters
1821
1822
1823
1824
1825
1826
1827
1828
1829
1830
1831
1832
1833
1834
1835
1836
1837
1838
1839
1840
1841
1842
1843
1844
1845
1846
1847
1848
1849
1850
1851
1852
1853
1854
1855
1856
1857
1858
1859
1860
1861
1862
1863
1864
1865
1866
1867
1868
1869
1870
1871
1872
1873
1874
1875
1876
1877
1878
1879
1880
1881
1882
1883
1884
1885
1886
1887
1888
1889
1890
1891
1892
1893
1894
1895
1896
1897
1898
1899
1900
1901
1902
1903
1904
1905
1906
1907
1908
1909
1910
1911
1912
1913
1914
1915
1916
1917
1918
1919
1920
1921
1922
1923
1924
1925
1926
1927
1928
1929
1930
1931
1932
1933
1934
1935
1936
1937
1938
1939
1940
1941
1942
1943
1944
1945
1946
1947
1948
1949
1950
1951
1952
1953
1954
1955
1956
1957
1958
1959
1960
1961
1962
1963
1964
1965
1966
1967
1968
1969
1970
1971
1972
1973
1974
1975
1976
1977
1978
1979
1980
1981
1982
1983
1984
1985
1986
1987
1988
1989
1990
1991
1992
1993
1994
1995
1996
1997
1998
1999
2000
2001
2002
2003
2004
2005
2006
2007
2008
2009
2010
2011
2012
2013
2014
2015
2016
2017
2018
2019
2020
2021
2022
2023
2024
2025
2026
2027
2028
2029
2030
2031
2032
2033
2034
2035
2036
2037
2038
2039
2040
2041
2042
2043
2044
2045
2046
2047
2048
2049
2050
2051
2052
2053
2054
2055
2056
2057
2058
2059
2060
2061
2062
2063
2064
2065
2066
2067
2068
2069
2070
2071
2072
2073
2074
2075
2076
2077
2078
2079
2080
2081
2082
2083
2084
2085
2086
2087
2088
2089
2090
2091
2092
2093
2094
2095
2096
2097
2098
2099
2100
2101
2102
2103
2104
2105
2106
2107
2108
2109
2110
2111
2112
2113
2114
2115
2116
2117
2118
2119
2120
2121
2122
2123
2124
2125
2126
2127
2128
2129
2130
2131
2132
2133
2134
2135
2136
2137
2138
2139
2140
2141
2142
2143
2144
2145
2146
2147
2148
2149
2150
2151
2152
2153
2154
2155
2156
2157
2158
2159
2160
2161
2162
2163
2164
2165
2166
2167
2168
2169
2170
2171
2172
2173
2174
2175
2176
2177
2178
2179
2180
2181
2182
2183
2184
2185
2186
2187
2188
2189
2190
2191
2192
2193
2194
2195
2196
2197
2198
2199
2200
2201
2202
2203
2204
2205
2206
2207
2208
2209
2210
2211
2212
2213
2214
2215
2216
2217
2218
2219
2220
2221
2222
2223
2224
2225
2226
2227
2228
2229
2230
2231
2232
2233
2234
2235
2236
2237
2238
2239
2240
2241
2242
2243
2244
2245
2246
2247
2248
2249
2250
2251
2252
2253
2254
2255
2256
2257
2258
2259
2260
2261
2262
2263
2264
2265
2266
2267
2268
2269
2270
2271
2272
2273
2274
2275
2276
2277
2278
2279
2280
2281
2282
2283
2284
2285
2286
2287
2288
2289
2290
2291
2292
2293
2294
2295
2296
2297
2298
2299
2300
2301
2302
2303
2304
2305
2306
2307
2308
2309
2310
2311
2312
2313
2314
2315
2316
2317
2318
2319
2320
2321
2322
2323
2324
2325
2326
2327
2328
2329
2330
2331
2332
2333
2334
2335
2336
2337
2338
2339
2340
2341
2342
2343
2344
2345
2346
2347
2348
2349
2350
2351
2352
2353
2354
2355
2356
2357
2358
2359
2360
2361
2362
2363
2364
2365
2366
2367
2368
2369
2370
2371
2372
2373
2374
2375
2376
2377
2378
2379
2380
2381
2382
2383
2384
2385
2386
2387
2388
2389
2390
2391
2392
2393
2394
2395
2396
2397
2398
2399
2400
2401
2402
2403
2404
2405
2406
2407
2408
2409
2410
2411
2412
2413
2414
2415
2416
2417
2418
2419
2420
2421
2422
2423
2424
2425
2426
2427
2428
2429
2430
2431
2432
2433
2434
2435
2436
2437
2438
2439
2440
2441
2442
2443
2444
2445
2446
2447
2448
2449
2450
2451
2452
2453
2454
2455
2456
2457
2458
2459
2460
2461
2462
2463
2464
2465
2466
2467
2468
2469
2470
2471
2472
2473
2474
2475
2476
2477
2478
2479
2480
2481
2482
2483
2484
2485
2486
2487
2488
2489
2490
2491
2492
2493
2494
2495
2496
2497
2498
2499
2500
2501
2502
2503
2504
2505
2506
2507
2508
2509
2510
2511
2512
2513
2514
2515
2516
2517
2518
2519
2520
2521
2522
2523
2524
2525
2526
2527
2528
2529
2530
2531
2532
2533
2534
2535
2536
2537
2538
2539
2540
2541
2542
2543
2544
2545
2546
2547
2548
2549
2550
2551
2552
2553
2554
2555
2556
2557
2558
2559
2560
2561
2562
2563
2564
2565
2566
2567
2568
2569
2570
2571
2572
2573
2574
2575
2576
2577
2578
2579
2580
2581
2582
2583
2584
2585
2586
2587
2588
2589
2590
2591
2592
2593
2594
2595
2596
2597
2598
2599
2600
2601
2602
2603
2604
2605
2606
2607
2608
2609
2610
2611
2612
2613
2614
2615
2616
2617
2618
2619
2620
2621
2622
2623
2624
2625
2626
2627
2628
2629
2630
2631
2632
2633
2634
2635
2636
2637
2638
2639
2640
2641
2642
2643
2644
2645
2646
2647
2648
2649
2650
2651
2652
2653
2654
2655
2656
2657
2658
2659
2660
2661
2662
2663
2664
2665
2666
2667
2668
2669
2670
2671
2672
2673
2674
2675
2676
2677
2678
2679
2680
2681
2682
2683
2684
2685
2686
2687
2688
2689
2690
2691
2692
2693
2694
2695
2696
2697
2698
2699
2700
2701
2702
2703
2704
2705
2706
2707
2708
2709
2710
2711
2712
2713
2714
2715
2716
2717
2718
2719
2720
2721
2722
2723
2724
2725
2726
2727
2728
2729
2730
2731
2732
2733
2734
2735
2736
2737
2738
2739
2740
2741
2742
2743
2744
2745
2746
2747
2748
2749
2750
2751
2752
2753
2754
2755
2756
2757
2758
2759
2760
2761
2762
2763
2764
2765
2766
2767
2768
2769
2770
2771
2772
2773
2774
2775
2776
2777
2778
2779
2780
2781
2782
2783
2784
2785
2786
2787
2788
2789
2790
2791
2792
2793
2794
2795
2796
2797
2798
2799
2800
2801
2802
2803
2804
2805
2806
2807
2808
2809
2810
2811
2812
2813
2814
2815
2816
2817
2818
2819
2820
2821
2822
2823
2824
2825
2826
2827
2828
2829
2830
2831
2832
2833
2834
2835
2836
2837
2838
2839
2840
2841
2842
2843
2844
2845
2846
2847
2848
2849
2850
2851
2852
2853
2854
2855
2856
2857
2858
2859
2860
2861
2862
2863
2864
2865
2866
2867
2868
2869
2870
2871
2872
2873
2874
2875
2876
2877
2878
2879
2880
2881
2882
2883
2884
2885
2886
2887
2888
2889
2890
2891
2892
2893
2894
2895
2896
2897
2898
2899
2900
2901
2902
2903
2904
2905
2906
2907
2908
2909
2910
2911
2912
2913
2914
2915
2916
2917
2918
2919
2920
2921
2922
2923
2924
2925
2926
2927
2928
2929
2930
2931
2932
2933
2934
2935
2936
2937
2938
2939
2940
2941
2942
2943
2944
2945
2946
2947
2948
2949
2950
2951
2952
2953
2954
2955
2956
2957
2958
2959
2960
2961
2962
2963
2964
2965
2966
2967
2968
2969
2970
2971
2972
2973
2974
2975
2976
2977
2978
2979
2980
2981
2982
2983
2984
2985
2986
2987
2988
2989
2990
2991
2992
2993
2994
2995
2996
2997
2998
2999
3000
3001
3002
3003
3004
3005
3006
3007
3008
3009
3010
3011
3012
3013
3014
3015
3016
3017
3018
3019
3020
3021
3022
3023
3024
3025
3026
3027
3028
3029
3030
3031
3032
3033
3034
3035
3036
3037
3038
3039
3040
3041
3042
3043
3044
3045
3046
3047
3048
3049
3050
3051
3052
3053
3054
3055
3056
3057
3058
3059
3060
3061
3062
3063
3064
3065
3066
3067
3068
3069
3070
3071
3072
3073
3074
3075
3076
3077
3078
3079
3080
3081
3082
3083
3084
3085
3086
3087
3088
3089
3090
3091
3092
3093
3094
3095
3096
3097
3098
3099
3100
3101
3102
3103
3104
3105
3106
3107
3108
3109
3110
3111
3112
3113
3114
3115
3116
3117
3118
3119
3120
3121
3122
3123
3124
3125
3126
3127
3128
3129
3130
3131
3132
3133
3134
3135
3136
3137
3138
3139
3140
3141
3142
3143
3144
3145
3146
3147
3148
3149
3150
3151
3152
3153
3154
3155
3156
3157
3158
3159
3160
3161
3162
3163
3164
3165
3166
3167
3168
3169
3170
3171
3172
3173
3174
3175
3176
3177
3178
3179
3180
3181
3182
3183
3184
3185
3186
3187
3188
3189
3190
3191
3192
3193
3194
3195
3196
3197
3198
3199
3200
3201
3202
3203
3204
3205
3206
3207
3208
3209
3210
3211
3212
3213
3214
3215
3216
3217
3218
3219
3220
3221
3222
3223
3224
3225
3226
3227
3228
3229
3230
3231
3232
3233
3234
3235
3236
3237
3238
3239
3240
3241
3242
3243
3244
3245
3246
3247
3248
3249
3250
3251
3252
3253
3254
3255
3256
3257
3258
3259
3260
3261
3262
3263
3264
3265
3266
3267
3268
3269
3270
3271
3272
3273
3274
3275
3276
3277
3278
3279
3280
3281
3282
3283
3284
3285
3286
3287
3288
3289
3290
3291
3292
3293
3294
3295
3296
3297
3298
3299
3300
3301
3302
3303
3304
3305
3306
3307
3308
3309
3310
3311
3312
3313
3314
3315
3316
3317
3318
3319
3320
3321
3322
3323
3324
3325
3326
3327
3328
3329
3330
3331
3332
3333
3334
3335
3336
3337
3338
3339
3340
3341
3342
3343
3344
3345
3346
3347
3348
3349
3350
3351
3352
3353
3354
3355
3356
3357
3358
3359
3360
3361
3362
3363
3364
3365
3366
3367
3368
3369
3370
3371
3372
3373
3374
3375
3376
3377
3378
3379
3380
3381
3382
3383
3384
3385
3386
3387
3388
3389
3390
3391
3392
3393
3394
3395
3396
3397
3398
3399
3400
3401
3402
3403
3404
3405
3406
3407
3408
3409
3410
3411
3412
3413
3414
3415
3416
3417
3418
3419
3420
3421
3422
3423
3424
3425
3426
3427
3428
3429
3430
3431
3432
3433
3434
3435
3436
3437
3438
3439
3440
3441
3442
3443
3444
3445
3446
3447
3448
3449
3450
3451
3452
3453
3454
3455
3456
3457
3458
3459
3460
3461
3462
3463
3464
3465
3466
3467
3468
3469
3470
3471
3472
3473
3474
3475
3476
3477
3478
3479
3480
3481
3482
3483
3484
3485
3486
3487
3488
3489
3490
3491
3492
3493
3494
3495
3496
3497
3498
3499
3500
3501
3502
3503
3504
3505
3506
3507
3508
3509
3510
3511
3512
3513
3514
3515
3516
3517
3518
3519
3520
3521
3522
3523
3524
3525
3526
3527
3528
3529
3530
3531
3532
3533
3534
3535
3536
3537
3538
3539
3540
3541
3542
3543
3544
3545
3546
3547
3548
3549
3550
3551
3552
3553
3554
3555
3556
3557
3558
3559
3560
3561
3562
3563
3564
3565
3566
3567
3568
3569
3570
3571
3572
3573
3574
3575
3576
3577
3578
3579
3580
3581
3582
3583
3584
3585
3586
3587
3588
3589
3590
3591
3592
3593
3594
3595
3596
3597
3598
3599
3600
3601
3602
3603
3604
3605
3606
3607
3608
3609
3610
3611
3612
3613
3614
3615
3616
3617
3618
3619
3620
3621
3622
3623
3624
3625
3626
3627
3628
3629
3630
3631
3632
3633
3634
3635
3636
3637
3638
3639
3640
3641
3642
3643
3644
3645
3646
3647
3648
3649
3650
3651
3652
3653
3654
3655
3656
3657
3658
3659
3660
3661
3662
3663
3664
3665
3666
3667
3668
3669
3670
3671
3672
3673
3674
3675
3676
3677
3678
3679
3680
3681
3682
3683
3684
3685
3686
3687
3688
3689
3690
3691
3692
3693
3694
3695
3696
3697
3698
3699
3700
3701
3702
3703
3704
3705
3706
3707
3708
3709
3710
3711
3712
3713
3714
3715
3716
3717
3718
3719
3720
3721
3722
3723
3724
3725
3726
3727
3728
3729
3730
3731
3732
3733
3734
3735
3736
3737
3738
3739
3740
3741
3742
3743
3744
3745
3746
3747
3748
3749
3750
3751
3752
3753
3754
3755
3756
3757

1836

Case 5: PPO (5/5)

1837

1838

Stage 7 - Final Result

1839

Final Answer

1840

The greatest length of rope that can be saved is $\boxed{160}$ meters.

1841

1842

Given five points $\langle A \rangle$, $\langle B \rangle$, $\langle C \rangle$, $\langle D \rangle$, and $\langle O \rangle$ on a flat field, where $\langle A \rangle$ is directly north of $\langle O \rangle$, $\langle B \rangle$ is directly west of $\langle O \rangle$, $\langle C \rangle$ is directly south of $\langle O \rangle$, and $\langle D \rangle$ is directly east of $\langle O \rangle$. The distance between $\langle C \rangle$ and $\langle D \rangle$ is 140 meters. A hot-air balloon is positioned at $\langle H \rangle$ directly above $\langle O \rangle$. The balloon is held by four ropes $\langle HA \rangle$, $\langle HB \rangle$, $\langle HC \rangle$, and $\langle HD \rangle$. The lengths of ropes $\langle HC \rangle$ and $\langle HD \rangle$ are 150 meters and 130 meters, respectively. We need to determine the greatest length of rope that can be saved by replacing $\langle HC \rangle$ and $\langle HD \rangle$ with a single rope $\langle HP \rangle$ where $\langle P \rangle$ is a point on the straight line between $\langle C \rangle$ and $\langle D \rangle$.

1843

1844

1845

1846

1847

First, we set up the coordinates for the points:

- $\langle O \rangle$ is at $\langle (0, 0) \rangle$.

- $\langle A \rangle$ is at $\langle (0, h) \rangle$, $\langle B \rangle$ is at $\langle (-k, 0) \rangle$, $\langle C \rangle$ is at $\langle (0, -m) \rangle$, and $\langle D \rangle$ is at $\langle (n, 0) \rangle$.

1848

Given:

- The distance between $\langle C \rangle$ and $\langle D \rangle$ is 140 meters, so $\sqrt{n^2 + m^2} = 140$.

- The distance $\langle HC \rangle$ is 150 meters, so $h + m = 150$.

- The distance $\langle HD \rangle$ is 130 meters, so $\sqrt{n^2 + h^2} = 130$.

1849

Solving these equations, we find:

- $h = 66$ meters, $m = 84$ meters, and $n = 112$ meters.

1850

1851

1852

1853

1854

To minimize the length of rope $\langle HP \rangle$, we parametrize point $\langle P \rangle$ on $\langle CD \rangle$ and find the point $\langle P \rangle$ that minimizes the distance $\langle HP \rangle$. Using calculus, we find the optimal parameter $t = \frac{9}{14}$, leading to the coordinates of $\langle P \rangle$ as $\langle (72, -30) \rangle$.

1855

1856

1857

1858

The minimal length of rope $\langle HP \rangle$ is calculated as 120 meters. The original total length of ropes $\langle HC \rangle$ and $\langle HD \rangle$ is 280 meters, so the greatest length of rope that can be saved is $280 - 120 = 160$ meters.

1859

1860

1861

1862

```
\[
\boxed{160}
]
```

<|end_of_sentence|>

1863

(3012 tokens)

1864

1865

1866

1867

1868

1869

1870

1871

1872

1873

1874

1875

1876

1877

1878

1879

1880

1881

1882

1883

1884

1885

1886

1887

1888

1889

1890

Case 2: Mix-PPG(1/4)

1891

1892

Stage 1 - Question Rehearsal

1893

Okay, so I have this problem with five points: A, B, C, D, and O on a flat field. A is north of O, B is west of O, C is south of O, and D is east of O. The distance between C and D is 140 meters. There's a hot-air balloon at point H directly above O. The balloon is held by four ropes: HA, HB, HC, and HD.

1894

Rope HC is 150 meters long, and HD is 130 meters long. The problem says we need to reduce the total length of rope by replacing HC and HD with a single rope HP, where P is somewhere on the line between C and D. I need to find the greatest length of rope that can be saved.

1895

Stage 2 - Spatial Configuration and Equation Derivation

1896

First, let me visualize this. Points A, B, C, D form a rectangle with O at the center. H is directly above O, so it's like a vertical pole. The ropes HA, HB, HC, HD are all connected to H and go to each of these points.

1897

Given that HC is 150 m and HD is 130 m. So the total length before replacement is $150 + 130 = 280$ m. After replacement, it's just one rope HP. So the maximum length we can save is $280 - \text{length of HP}$. So I need to find the minimum possible length of HP, which would give the maximum saving.

1898

So, to minimize HP, point P should be chosen such that the distance from H to P is as small as possible. Since P is on CD, which is a straight line between C and D, which are 140 meters apart.

1899

Let me assign coordinates to make this easier. Let's place O at $(0, 0)$. Then, since A is north of O, A is $(0, a)$ for some a. Similarly, B is west of O, so B is $(-b, 0)$ for some b. C is south of O, so C is $(0, -c)$, and D is east of O, so D is $(d, 0)$.

1900

Stage 3 - Step-by-Step Computation

1901

Given that the distance between C and D is 140 m. Since C is $(0, -c)$ and D is $(d, 0)$, the distance between them is $\sqrt{(d - 0)^2 + (0 - (-c))^2} = \sqrt{d^2 + c^2} = 140$. So $d^2 + c^2 = 140^2 = 19600$.

1902

Now, the balloon is at H, which is directly above O, so H is $(0, h)$ for some h.

1903

Given that HC is 150 m. So the distance from H to C is $\sqrt{(0 - 0)^2 + (-c - h)^2} = | -c - h | = |c + h| = 150$. So $c + h = 150$. Similarly, HD is 130 m. So the distance from H to D is $\sqrt{(d - 0)^2 + (0 - h)^2} = \sqrt{d^2 + h^2} = 130$. So $d^2 + h^2 = 130^2 = 16900$.

1904

So now, we have two equations:

1905

1. $c + h = 150$
2. $d^2 + h^2 = 16900$
3. $c^2 + d^2 = 19600$

1906

From equation 1, $c = 150 - h$. Plug this into equation 3:

1907

$$(150 - h)^2 + d^2 = 19600$$

1908

Expand $(150 - h)^2$: $22500 - 300h + h^2$

1909

$$\text{So } 22500 - 300h + h^2 + d^2 = 19600$$

1910

1911

1912

1913

1914

1915

1916

1917

1918

1919

1920

1921

1922

1923

1924

1925

1926

1927

1928

1929

1930

1931

1932

1933

1934

1935

1936

1937

1938

1939

1940

1941

1942

1943

1944 Case 2: Mix-PPG(2/4)

1945 But from equation 2, $d^2 = 16900 - h^2$. Substitute that in:

1946 $22500 - 300h + h^2 + (16900 - h^2) = 19600$

1947 Simplify:

1948 $22500 - 300h + h^2 + 16900 - h^2 = 19600$

1949 The h^2 terms cancel:

1950 $22500 + 16900 - 300h = 19600$

1951 Compute $22500 + 16900 = 39400$

1952 So $39400 - 300h = 19600$

1953 Subtract 19600 from both sides:

1954 $39400 - 19600 - 300h = 0$

1955 $19800 - 300h = 0$

1956 So $300h = 19800$

1957 $h = 19800 / 300 = 66$

1958 So h is 66 meters. Then $c = 150 - h = 150 - 66 = 84$ meters.

1959 Then from equation 2, $d^2 + h^2 = 16900$

1960 $d^2 + 66^2 = 16900$

1961 66^2 is 4356, so $d^2 = 16900 - 4356 = 12544$

1962 So $d = \sqrt{12544} = 112$ meters.

1963 So now, we have $h = 66$, $c = 84$, $d = 112$.

1964 **## Stage 4 - Optimization Strategy**

1965 Now, point P is somewhere on CD. Let's parameterize point P. Since CD is from C (0, -84) to D (112, 0). So the line CD can be parameterized as $(112t, -84(1 - t))$ where t is from 0 to 1.

1966 So point P is $(112t, -84(1 - t))$.

1967 We need to find t such that the distance from H (0, 66) to P $(112t, -84(1 - t))$ is minimized.

1968 **## Stage 5 - Step-by-Step Computation**

1969 Compute the distance squared:

1970 $(112t - 0)^2 + (-84(1 - t) - 66)^2$

1971 $= (112t)^2 + (-84(1 - t) - 66)^2$

1972

1973

1974

1975

1976

1977

1978

1979

1980

1981

1982

1983

1984

1985

1986

1987

1988

1989

1990

1991

1992

1993

1994

1995

1996

1997

1998 **Case 2: Mix-PPG (3/4)**

1999

2000 Compute each part:

2001 First term: $(112t)^2 = 12544 t^2$

2002 Second term: $(-84 + 84t - 66) = (-150 + 84t)$. So squared: $(-150 + 84t)^2$

2003 Compute $(-150 + 84t)^2: 150^2 - 2*150*84t + (84t)^2 = 22500 - 25200t + 7056 t^2$

2004 So total distance squared is:

2005 $12544 t^2 + 22500 - 25200t + 7056 t^2$

2006 Combine like terms:

2007 $(12544 + 7056) t^2 - 25200t + 22500$

2008 Compute $12544 + 7056: 19600$

2009 So distance squared is $19600 t^2 - 25200t + 22500$

2010 To find the minimum, take derivative with respect to t and set to zero.

2011 Let $f(t) = 19600 t^2 - 25200t + 22500$

2012 $f'(t) = 39200 t - 25200$

2013 Set to zero:

2014 $39200 t - 25200 = 0$

2015 $39200 t = 25200$

2016 $t = 25200 / 39200$

2017 Simplify: divide numerator and denominator by 2800: $25200 / 2800 = 9$, $39200 / 2800 = 14$

2018 So $t = 9/14$

2019 So t is $9/14$. So point P is at $(112*(9/14), -84*(1 - 9/14))$.

2020 Compute $112*(9/14)$: 112 divided by 14 is 8, so $8*9=72$. So x-coordinate is 72.

2021 Compute $-84*(1 - 9/14)$: $1 - 9/14 = 5/14$. So $-84*(5/14) = -6*5 = -30$. So y-coordinate is -30.

2022 So point P is $(72, -30)$.

2023 Now, compute the distance HP.

2024 Coordinates of H: $(0, 66)$. Coordinates of P: $(72, -30)$.

2025 Distance squared: $(72 - 0)^2 + (-30 - 66)^2 = 72^2 + (-96)^2$

2026 Compute $72^2: 5184$

2027 Compute $(-96)^2: 9216$

2028

2029

2030

2031

2032

2033

2034

2035

2036

2037

2038

2039

2040

2041

2042

2043

2044

2045

2046

2047

2048

2049

2050

2051

2052

Case 2: Mix-PPG(4/4)

2053

2054

Total distance squared: $5184 + 9216 = 14400$

2055

So distance is $\sqrt{14400} = 120$ meters.

2056

2057

So the length of HP is 120 meters. Therefore, the total length saved is $280 - 120 = 160$ meters.

2058

Stage 6 - Final Answer

2059

Final Answer

2060

The greatest length of rope that can be saved is $\boxed{160}$ meters.

2061

</think>

Given five points $\langle A \rangle$, $\langle B \rangle$, $\langle C \rangle$, $\langle D \rangle$, and $\langle O \rangle$ on a flat field, where $\langle A \rangle$ is directly north of $\langle O \rangle$, $\langle B \rangle$ is directly west of $\langle O \rangle$, $\langle C \rangle$ is directly south of $\langle O \rangle$, and $\langle D \rangle$ is directly east of $\langle O \rangle$. The distance between $\langle C \rangle$ and $\langle D \rangle$ is 140 meters. A hot-air balloon is positioned at $\langle H \rangle$ directly above $\langle O \rangle$. The balloon is held by four ropes $\langle HA \rangle$, $\langle HB \rangle$, $\langle HC \rangle$, and $\langle HD \rangle$. Rope $\langle HC \rangle$ is 150 meters and rope $\langle HD \rangle$ is 130 meters. We need to determine the greatest length of rope that can be saved by replacing $\langle HC \rangle$ and $\langle HD \rangle$ with a single rope $\langle HP \rangle$ where $\langle P \rangle$ is a point on the straight line between $\langle C \rangle$ and $\langle D \rangle$.

2062

First, we assign coordinates to the points. Let $\langle O \rangle$ be at $\langle (0, 0) \rangle$. Points $\langle A \rangle$, $\langle B \rangle$, $\langle C \rangle$, and $\langle D \rangle$ are at $\langle (0, a) \rangle$, $\langle (-b, 0) \rangle$, $\langle (0, -c) \rangle$, and $\langle (d, 0) \rangle$ respectively. The distance between $\langle C \rangle$ and $\langle D \rangle$ is 140 meters, giving us the equation $\langle d^2 + c^2 = 19600 \rangle$.

2063

Given the distances $\langle HC = 150 \rangle$ meters and $\langle HD = 130 \rangle$ meters, we have:

2064

1. $\langle c + h = 150 \rangle$

2065

2. $\langle d^2 + h^2 = 16900 \rangle$

2066

3. $\langle c^2 + d^2 = 19600 \rangle$

2067

Solving these equations, we find:

2068

- $\langle c = 84 \rangle$ meters

2069

- $\langle h = 66 \rangle$ meters

2070

- $\langle d = 112 \rangle$ meters

2071

Point $\langle P \rangle$ on $\langle CD \rangle$ is parameterized as $\langle (112t, -84(1 - t)) \rangle$. The distance from $\langle H \rangle$ to $\langle P \rangle$ is minimized by finding the value of $\langle t \rangle$ that minimizes the distance squared function. The minimum occurs at $\langle t = \frac{9}{14} \rangle$, giving point $\langle P \rangle$ as $\langle (72, -30) \rangle$.

2072

The distance $\langle HP \rangle$ is calculated as:

2073

- Distance squared: $\langle 72^2 + (-96)^2 = 14400 \rangle$

2074

- Distance: $\langle \sqrt{14400} = 120 \rangle$ meters

2075

Thus, the greatest length of rope that can be saved is $\langle (280 - 120 = 160) \rangle$ meters.

2076

The final answer is $\langle \boxed{160} \rangle$.<|end_of_sentence|>

(2830 tokens)

2077

2078

2079

2080

2081

2082

2083

2084

2085

2086

2087

2088

2089

2090

2091

2092

2093

2094

2095

2096

2097

2098

2099

2100

2101

2102

2103

2104

2105

2106

Case 3: Mix-PPG with an Increased UTD Ratio

2107

Stage 1 - Question Rehearsal

2109

First, I identify the initial lengths of ropes HC and HD. Both are given as 150 m and 130 m respectively, totaling 280 m.

2110

Next, I analyze the new rope HP. Point P lies somewhere on the line segment CD. Since C and D are directly south and north of O respectively, CD is a vertical line. The balloon is positioned directly above O, so the horizontal distance from O to H is the same as the height of the balloon.

2111

Stage 2 - Spatial Configuration

2112

Using the Pythagorean theorem for triangle HCP, where HC is 150 m and CP is the variable distance from C to P, I set up the equation:

2113

$$\text{Sqrt}(140^2 + h^2) = 150 \text{ m.}$$

2114

Solving for h, the height of the balloon, I find $h = \text{Sqrt}(150^2 - 140^2) = \text{Sqrt}(22100) \approx 148.66 \text{ m.}$

2115

Then, for triangle HDP, with DP = CD - CP = 140 - x, I set up the equation:

$$\text{Sqrt}(130^2 + h^2) = \text{Sqrt}(130^2 + 22100).$$

2116

Calculating this gives the new length of HP as $\text{sqrt}(16900 + 22100) = \text{sqrt}(38200) \approx 195.44 \text{ m.}$

2117

Adding both new lengths gives approximately 390.88 m, compared to the original 280 m. The difference is approximately 110.88 m, which is the maximum length of rope that can be saved.

</think>

2118

Stage 3 - Step-by-Step Computation

2119

To determine the maximum length of rope that can be saved by replacing ropes HC and HD with a single rope HP, we analyze the situation step by step.

2120

****Initial Setup:****

- ****Ropes HC and HD:**** Given lengths are 150 m and 130 m, respectively.
- ****Total Initial Length:**** $\text{sqrt}(150^2 + 130^2) = 280 \text{ m.}$

2121

****After Replacement:****

- ****New Rope HP:**** Point P lies on segment CD, replacing HC and HD.
- ****Coordinates and Distances:****
- Assume O is at the origin, H is vertically above O at (0, h).
- C and D are vertically aligned with C below and D above O.
- Using Pythagoras for triangles HCP and HDP:

1. For HC: $\text{sqrt}(140^2 + h^2) = 150 \rightarrow h = \text{sqrt}(150^2 - 140^2) \approx 148.66 \text{ m.}$

2. For DP: $\text{sqrt}(130^2 + h^2) = \text{sqrt}(130^2 + 148.66^2) \approx 195.44 \text{ m.}$

2122

****Calculation of Savings:****

- ****Total New Length:**** $\text{sqrt}(150^2 + 195.44^2) \approx 349.44 \text{ m.}$
- ****Savings:**** $280 - 349.44 \approx 68.56 \text{ m.}$

2123

However, precise calculation shows maximum saving occurs when P is optimally placed, yielding exact $60\sqrt{10} \text{ m} \approx 187.08 \text{ m total, saving } \approx 92.92 \text{ m.}$ But standard approach gives approximate 110.88 m saving.

2124

Stage 4 - Final Answer

2125

****Final Answer:****

```
\[
\boxed{110.88}
\] <|end_of_sentence|>
```

(908 tokens)

2126

2127

2128

2129

2130

2131

2132

2133

2134

2135

2136

2137

2138

2139

2140

2141

2142

2143

2144

2145

2146

2147

2148

2149

2150

2151

2152

2153

2154

2155

2156

2157

2158

2159