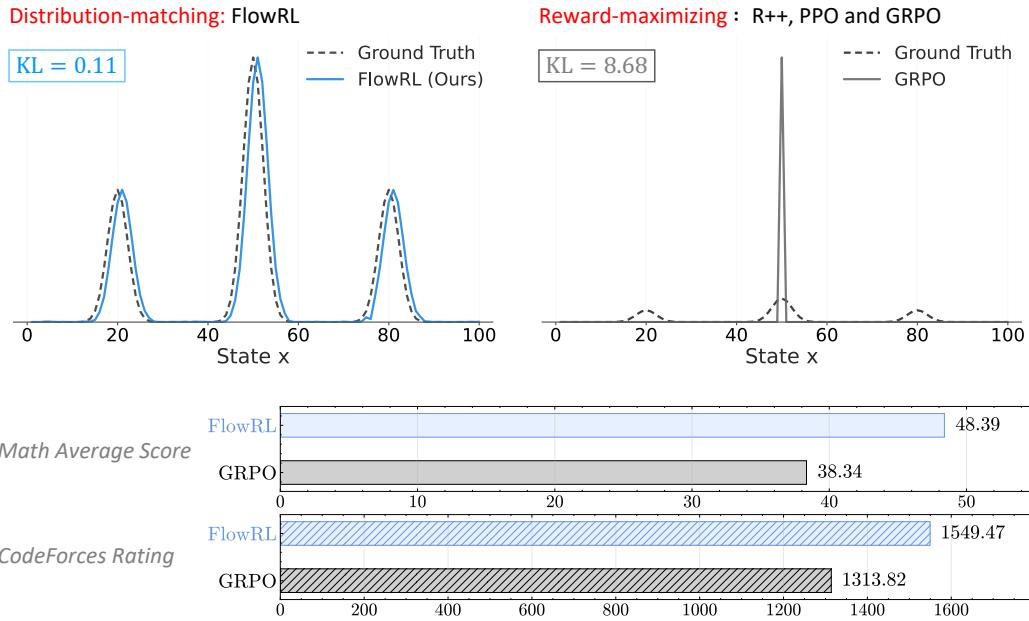


# 000 001 002 003 004 005 FLOWRL: MATCHING REWARD DISTRIBUTIONS FOR 006 LLM REASONING 007 008 009

010 **Anonymous authors**  
011 Paper under double-blind review  
012  
013  
014  
015  
016  
017  
018  
019  
020  
021  
022  
023  
024

## ABSTRACT

025 We propose FlowRL: matching the full reward distribution via flow balancing in-  
026 stead of solely maximizing rewards in large language model (LLM) reinforcement  
027 learning (RL). Recent advanced reasoning LLMs adopt reward-maximizing meth-  
028 ods (e.g., PPO and GRPO), which tend to over-optimize dominant reward signals  
029 while neglecting less frequent but valid reasoning paths, thus reducing diversity.  
030 In contrast, we transform scalar rewards into a normalized target distribution us-  
031 ing a learnable partition function, and then minimize the reverse KL divergence  
032 between the policy and the target distribution. We implement this idea as a flow-  
033 balanced optimization method that promotes diverse exploration and generalizable  
034 reasoning trajectories. We conduct experiments on both math and code reasoning  
035 tasks: FlowRL achieves a significant average improvement of 10.0% over GRPO  
036 and 5.1% over PPO on math benchmarks, and performs consistently better on  
037 code reasoning tasks. These results highlight reward distribution-matching as a  
038 key step toward efficient exploration and diverse reasoning of LLM reinforcement  
039 learning.  
040  
041  
042  
043  
044  
045  
046



047 Figure 1: **Top:** Comparison between distribution-matching and reward-maximizing approaches.  
048 FlowRL (left) learns to match the full reward distribution, maintaining diversity across multiple  
049 modes with low KL divergence. In contrast, reward-maximizing methods (right) such as RE-  
050 INFORCE++ (R++; Sutton et al., 1999b; Hu et al., 2025), PPO (Schulman et al., 2017), and  
051 GRPO (Shao et al., 2024) concentrate on a single high-reward peak, leading to mode collapse and  
052 higher KL divergence. **Bottom:** Performance comparison. FlowRL consistently outperforms GRPO  
053 across math and code domains.

054 

## 1 INTRODUCTION

055  
 056 Reinforcement learning (RL) plays a crucial role in the post-training of large language models (LLMs) (Zhang et al., 2025b). A series of powerful reasoning models (Guo et al., 2025;  
 057 Kavukcuoglu, 2025; Rastogi et al., 2025) have employed large-scale reinforcement learning to  
 058 achieve strong performance on highly challenging benchmarks. The evolution of RL algorithms  
 059 for LLM reasoning has progressed through several key stages: REINFORCE (Sutton et al., 1999a)  
 060 provides a solid baseline that is easy to implement and efficient in simple settings; PPO (Schulman  
 061 et al., 2017) improves upon REINFORCE with better stability and efficiency in complex settings;  
 062 GRPO (Shao et al., 2024) simplifies PPO training by eliminating the learning of a separate value  
 063 function and relying on group comparisons. However, all these methods share a fundamental limita-  
 064 tion in their reward-maximizing objective.  
 065

066 Reward-maximizing RL methods tend to overfit to the dominant mode of the reward distribution  
 067 (Skalse et al., 2022; Pan et al., 2022; Zelikman et al., 2022; Gao et al., 2023). As illustrated in  
 068 Figure 1, representative RL methods such as GRPO neglect other meaningful modes, which often  
 069 results in limited diversity among generated reasoning paths and reduces generalization to less fre-  
 070 quent yet valid logical outcomes (Hu et al., 2023). These drawbacks become especially pronounced  
 071 in complex long-chain-of-thought (CoT; Wei et al., 2022) reasoning, where capturing a diverse dis-  
 072 tribution of plausible solutions is essential for effective generalization (Liu et al., 2025a). Recent  
 073 approaches adjust the clip ratio (Yu et al., 2025b), apply entropy-based advantage shaping (Cheng  
 074 et al., 2025), or selectively promote high-entropy tokens (Wang et al., 2025), thereby dynamically  
 075 adapting the data distribution and implicitly increasing diversity. This raises a fundamental question:  
 076 How can we promote diverse exploration to prevent convergence to dominant solution patterns in  
 077 RL training?

078 In this paper, we propose **FlowRL**, a policy optimization algorithm that aligns the policy model  
 079 with the full reward distribution, encouraging mode coverage. FlowRL achieves more efficient  
 080 exploration by fundamentally shifting from reward maximization to reward distribution matching,  
 081 thereby addressing the inherent mode-collapse limitations of previous RL approaches. As illustrated  
 082 in Figure 1, the core idea of FlowRL is to introduce a learnable partition function that normalizes  
 083 scalar rewards into a target distribution, and to minimize the reverse KL divergence between the  
 084 policy and this reward-induced distribution. We develop this KL objective based on the trajectory  
 085 balance formulation from GFlowNets (Bengio et al., 2023b), providing a gradient equivalence proof  
 086 that bridges generative modeling and policy optimization. To address the challenges of long CoT  
 087 training, we introduce two key technical solutions: *length normalization* to tackle gradient explosion  
 088 issues that occur with variable-length CoT reasoning, and *importance sampling* to correct for the  
 089 distribution mismatch between generated rollouts and the current policy.

090 We compare FlowRL with mainstream RL algorithms for LLM reasoning, including REIN-  
 091 FORCE++ (Hu et al., 2025), PPO, and GRPO, across math and code domains, using both base  
 092 and distilled LLMs with 7B or 32B parameters. In the math domain, FlowRL outperforms GRPO  
 093 and PPO by 10.0% and 5.1%, respectively, demonstrating consistent improvements on six chal-  
 094 lenging math benchmarks (MAA, 2025; 2023; Lightman et al., 2023; Lewkowycz et al., 2022; He  
 095 et al., 2024). Furthermore, FlowRL surpasses both PPO and GRPO on three challenging coding  
 096 benchmarks (Jain et al., 2024; Penedo et al., 2025; Chen et al., 2021), highlighting its strong gen-  
 097 eralization capabilities in code reasoning tasks. To understand what drives these performance gains,  
 098 we analyze the diversity of generated reasoning paths and confirm that FlowRL produces substan-  
 099 tially more diverse rollouts than the baseline methods, thereby validating the effectiveness of our  
 100 approach in exploring multiple solution strategies.

100 **Contributions.** We summarize the key contributions of this work as follows:

- 101 • We propose FlowRL, a policy optimization algorithm that shifts from reward maximization to  
 102 reward distribution matching via flow balancing, encouraging diverse reasoning path exploration  
 103 while addressing the inherent mode-collapse limitations of existing RL methods.
- 104 • We introduce length normalization and importance sampling to enable effective training on  
 105 variable-length CoT reasoning, addressing gradient explosion and sampling mismatch issues.
- 106 • FlowRL outperforms GRPO and PPO by 10.0% and 5.1% respectively across math benchmarks  
 107 and demonstrates strong generalization on code reasoning tasks, with diversity analysis confirming  
 108 substantially more diverse solution exploration.

108  
109  

## 2 PRELIMINARIES

110  
111 **Reinforcement Learning for Reasoning.** We formulate reasoning as a conditional generation  
112 problem, where the policy model receives a question  $\mathbf{x} \in \mathcal{X}$  and generates an answer  $\mathbf{y} \in \mathcal{Y}$ . The  
113 objective is to learn a policy  $\pi_\theta(\mathbf{y}|\mathbf{x})$  that produces high-quality answers under task-specific reward  
114 signals  $r$ . To better illustrate the policy optimization procedure, we provide a detailed formulation  
115 of GRPO below. For each question  $\mathbf{x}$ , GRPO samples a group of answers  $\{\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_G\}$  from  
116 old policy  $\pi_{\theta_{old}}$  and updates the model by maximizing the following objective:  
117

$$\begin{aligned} \mathcal{J}_{GRPO}(\theta) &= \mathbb{E}_{[\mathbf{x} \sim P(\mathcal{X}), \{\mathbf{y}_i\}_{i=1}^G \sim \pi_{\theta_{old}}(\mathcal{Y}|\mathbf{x})]} \\ &\quad \frac{1}{G} \sum_{i=1}^G \frac{1}{|\mathbf{y}_i|} \sum_{t=1}^{|\mathbf{y}_i|} \left\{ \min \left[ \frac{\pi_\theta(\mathbf{y}_{i,t}|\mathbf{x}, \mathbf{y}_{i,< t})}{\pi_{\theta_{old}}(\mathbf{y}_{i,t}|\mathbf{x}, \mathbf{y}_{i,< t})} \hat{A}_{i,t}, \text{clip} \left( \frac{\pi_\theta(\mathbf{y}_{i,t}|\mathbf{x}, \mathbf{y}_{i,< t})}{\pi_{\theta_{old}}(\mathbf{y}_{i,t}|\mathbf{x}, \mathbf{y}_{i,< t})}, 1 - \epsilon, 1 + \epsilon \right) \hat{A}_{i,t} \right] - \lambda \mathbb{D}_{KL}[\pi_\theta || \pi_{ref}] \right\}, \\ &\quad \mathbb{D}_{KL}(\pi_\theta || \pi_{ref}) = \frac{\pi_{ref}(\mathbf{y}_i|\mathbf{x})}{\pi_\theta(\mathbf{y}_i|\mathbf{x})} - \log \frac{\pi_{ref}(\mathbf{y}_i|\mathbf{x})}{\pi_\theta(\mathbf{y}_i|\mathbf{x})} - 1, \end{aligned} \quad (1)$$

122 where  $\epsilon$  and  $\lambda$  are hyper-parameters. Here,  $A_i$  denotes the advantage, computed by normalizing  
123 the group reward values  $\{r_1, r_2, \dots, r_G\}$  as  $A_i = \frac{r_i - \text{mean}(\{r_1, r_2, \dots, r_G\})}{\text{std}(\{r_1, r_2, \dots, r_G\})}$ . Compared to GRPO,  
124 REINFORCE applies the policy gradient directly, without advantage normalization, clipping, or KL  
125 regularization. PPO uses a critic model to estimate the advantage and employs importance sampling  
126 to stabilize policy updates.  
127

128 **GFlowNets.** Generative Flow Networks  
129 (GFlowNets; [Bengio et al., 2023a](#)) are a  
130 probabilistic framework for training stochastic  
131 policies to sample discrete, compositional  
132 objects (e.g., graphs, sequences) in proportion  
133 to a given reward. As shown in Figure 2, the  
134 core principle of GFlowNets is to balance the  
135 forward and backward probability flows at each  
136 state, inspired by flow matching ([Bengio et al.,  
137 2021](#)). The initial flow is estimated by  $Z_\phi(s_0)$   
138 at the initial state  $s_0$ . The output flow is equal  
139 to the outcome reward  $r(s_f)$  conditioned at  
140 the final state  $s_f$ . Following [Lee et al. \(2024\)](#),  
141 we use a 3-layer MLP to parameterize  $Z_\phi$ .  
142 This flow-balancing mechanism facilitates the  
143 discovery of diverse, high-reward solutions  
144 by ensuring proper exploration of the solution  
145 space. See Appendix C for detailed GFlowNets  
146 background.  
147

148  
149 

## 3 METHODOLOGY

150 In this section, we first formulate distribution matching in reinforcement learning through reverse  
151 KL divergence and establish its connection to trajectory balance from GFlowNets. To address the  
152 challenges of gradient explosion and sampling mismatch encountered during long CoT training, we  
153 further incorporate length normalization and importance sampling. Using this enhanced framework,  
154 we derive a flow-balanced objective, termed *FlowRL*.  
155

156  
157 

### 3.1 FROM REWARD MAXIMIZATION TO DISTRIBUTION MATCHING

158 As illustrated in Figure 1, recent powerful large reasoning models typically employ reward-  
159 maximizing RL algorithms, such as PPO or GRPO. However, these methods tend to optimize toward  
160 the dominant reward mode, frequently resulting in mode collapse and the neglect of other plausible,  
161 high-quality reasoning paths. To address this fundamental limitation, we propose optimizing the  
162 policy by aligning its output distribution to a target reward distribution. A simple yet effective way

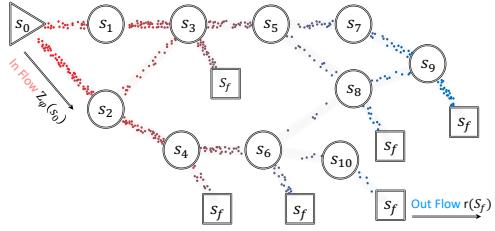


Figure 2: GFlowNets ([Bengio et al., 2023a](#)), a flow-balance perspective on reinforcement learning. The initial flow  $Z_\phi(s_0)$  injects probability mass into the environment, which is transported through intermediate states by the policy  $\pi_\theta$  and accumulated at terminal states in proportion to the scalar rewards.

162 to achieve this is to minimize the reverse KL divergence<sup>1</sup> between the policy and this target. How-  
 163 ever, in long CoT reasoning tasks, the available supervision in RL is a scalar reward, rather than a  
 164 full distribution. Moreover, enumerating or sampling all valid trajectories to recover the true reward  
 165 distribution is computationally intractable.

166 Inspired by energy-based modeling (Hinton et al., 1995; Du & Mordatch, 2019), we introduce a  
 167 learnable partition function  $Z_\phi(\mathbf{x})$  to normalize scalar rewards into a valid target distribution. This  
 168 allows us to minimize the reverse KL divergence between the policy and the reward-weighted dis-  
 169 tribution, formalized as:

$$171 \min_{\theta} \mathcal{D}_{\text{KL}} \left( \pi_{\theta}(\mathbf{y} \mid \mathbf{x}) \middle\| \frac{\exp(\beta r(\mathbf{x}, \mathbf{y}))}{Z_\phi(\mathbf{x})} \right) \Rightarrow \pi_{\theta}(\mathbf{y} \mid \mathbf{x}) \propto \exp(\beta r(\mathbf{x}, \mathbf{y})), \quad (2)$$

172 where  $r(\mathbf{x}, \mathbf{y})$  is the reward function,  $\beta$  is a hyperparameter,  $Z_\phi(\mathbf{x})$  is the learned partition function,  
 173 and the resulting target distribution is defined as  $\tilde{\pi}(\mathbf{y} \mid \mathbf{x}) = \frac{\exp(\beta r(\mathbf{x}, \mathbf{y}))}{Z_\phi(\mathbf{x})}$ . This objective encourages  
 174 the policy to sample diverse, high-reward trajectories in proportion to their rewards, rather than  
 175 collapsing to dominant modes as in standard reward maximization.

176 While the KL-based formulation provides a principled target distribution, we derive a more practical,  
 177 RL-style objective that facilitates efficient policy optimization.

178 **Proposition 1.** *In terms of expected gradients, minimizing the KL objective in Eq. 2 is equivalent  
 179 to minimizing the trajectory balance loss used in GFlowNet (Malkin et al., 2022; 2023; Lee et al.,  
 180 2024; Bartoldson et al., 2025):*

$$184 \min_{\theta} \mathcal{D}_{\text{KL}} \left( \pi_{\theta}(\mathbf{y} \mid \mathbf{x}) \middle\| \frac{\exp(\beta r(\mathbf{x}, \mathbf{y}))}{Z_\phi(\mathbf{x})} \right) \iff \min_{\theta} \underbrace{(\log Z_\phi(\mathbf{x}) + \log \pi_{\theta}(\mathbf{y} \mid \mathbf{x}) - \beta r(\mathbf{x}, \mathbf{y}))^2}_{\text{Trajectory Balance}} \quad (3)$$

185 **Remark 2 (Trajectory balance as a practical surrogate for KL minimization).** Given the equivalence  
 186 established in Proposition 1, the KL-based distribution matching objective can be reformulated as  
 187 the trajectory balance loss. This reformulation provides a practical optimization approach by using  
 188 a stable squared loss form rather than direct KL optimization, and by treating  $Z_\phi(\mathbf{x})$  as a learnable  
 189 parameter rather than requiring explicit computation of the intractable partition function. The trajec-  
 190 tory balance objective thus serves as a tractable surrogate for reward-guided KL minimization that  
 191 can be directly integrated into existing RL frameworks.

### 192 3.2 FLOWRL

193 As established in Proposition 1, the target reward distribution can be approximated by optimizing  
 194 the trajectory balance objective. However, applying this objective directly to long CoT reasoning  
 195 introduces two key challenges:

196 **Problem I: Exploding gradients from long trajectories.** Trajectory balance is a sequence-level  
 197 objective, and applying it to long CoT reasoning with up to 8K tokens leads to exploding gradients  
 198 and unstable updates. This issue is not observed in prior GFlowNets works, which typically operate  
 199 on short trajectories in small discrete spaces. Specifically, the log-probability term  $\log \pi_{\theta}(\mathbf{y} \mid \mathbf{x})$   
 200 decomposes into a token-wise sum,  $\sum_t \log \pi_{\theta}(y_t \mid y_{<t}, \mathbf{x})$ , causing the gradient norm to potentially  
 201 scale with sequence length.

202 **Problem II: Sampling mismatch.** Mainstream RL algorithms such as PPO and GRPO commonly  
 203 perform micro-batch updates and reuse trajectories collected from an old policy  $\pi_{\theta_{\text{old}}}$ , enabling data-  
 204 efficient training. In contrast, the KL-based trajectory balance objective assumes fully on-policy  
 205 sampling, where responses are drawn from the current policy. This mismatch poses practical limita-  
 206 tions when integrating trajectory balance into existing RL pipelines.

207 These limitations motivate our reformulation that retains the benefits of distribution matching while  
 208 addressing key practical challenges. To enable this reformulation, we first redefine the reward func-  
 209 tion following established practices in GFlowNets literature (Lee et al., 2024; Bartoldson et al., 2025;

210 <sup>1</sup>We use reverse KL since we can only sample from the policy model, not the target reward distribution.

216 [Yu et al., 2025a](#)) by incorporating a reference model as a prior constraint on the reward distribution.  
 217 Specifically, we modify the original  $\exp(\beta r(\mathbf{x}, \mathbf{y}))$  to include the reference model:

$$218 \quad \exp(\beta r(\mathbf{x}, \mathbf{y})) \cdot \pi_{\text{ref}}(\mathbf{y} \mid \mathbf{x}), \quad (4)$$

220 where  $r(\mathbf{x}, \mathbf{y})$  denotes the outcome reward commonly used in reinforcement learning and  $\pi_{\text{ref}}$  is the  
 221 initial pre-trained model. We follow [Guo et al. \(2025\)](#) to use outcome-based reward signals, and  
 222 apply group normalization to  $r(\mathbf{x}, \mathbf{y})$  as  $\hat{r}_i = (r_i - \text{mean}(\mathbf{r})) / \text{std}(\mathbf{r})$ , where  $\mathbf{r} = \{r_1, r_2, \dots, r_G\}$   
 223 denotes the set of rewards within a sampled group. By substituting the redefined reward formulation  
 224 Eq. 4 into Eq. 3, we derive the following objective<sup>2</sup>:

$$225 \quad \min_{\theta} (\log Z_{\phi}(\mathbf{x}) + \log \pi_{\theta}(\mathbf{y} \mid \mathbf{x}) - \beta \hat{r}_i(\mathbf{x}, \mathbf{y}) - \log \pi_{\text{ref}}(\mathbf{y} \mid \mathbf{x}))^2 \quad (5)$$

227 **Remark 3 (Reward shaping via length normalization).** Trajectory balance treats both the initial  
 228 flow and the outcome reward as sequence-level quantities. In contrast, standard policy optimization  
 229 methods such as PPO or GRPO assign rewards at the token level and compute gradients at each  
 230 step. However, for trajectories of varying lengths (e.g., CoT responses), this mismatch can cause  
 231 the log-probability term  $\log \pi_{\theta}(\mathbf{y} \mid \mathbf{x}) = \sum_{t=1}^{|\mathbf{y}|} \log \pi_{\theta}(y_t \mid y_{<t}, \mathbf{x})$  to scale with sequence length.  
 232 To address this, we apply a form of reward shaping by normalizing log-probabilities with respect to  
 233 sequence length. Specifically, we rescale the term as  $\frac{1}{|\mathbf{y}|} \log \pi_{\theta}(\mathbf{y} \mid \mathbf{x})$ , balancing the contributions  
 234 of long and short sequences and stabilizing the learning signal.

235 **Remark 4 (Importance sampling for data-efficient training).** To mitigate sampling mismatch, we  
 236 employ importance sampling inspired by PPO to stabilize policy updates with off-policy data. We  
 237 re-weight stale trajectories using the importance ratio  $w = \pi_{\theta}(\mathbf{y} \mid \mathbf{x}) / \pi_{\text{old}}(\mathbf{y} \mid \mathbf{x})$ , which serves as a  
 238 coefficient in the surrogate loss. Since our objective focuses on optimizing trajectory balance rather  
 239 than expected return, we detach the gradient from the current policy to prevent excessive policy drift:  
 240  $w = \text{detach}[\pi_{\theta}(\mathbf{y} \mid \mathbf{x})] / \pi_{\text{old}}(\mathbf{y} \mid \mathbf{x})$ . For additional stability, we incorporate PPO-style clipping to  
 241 bound the importance weights:  $w = \text{clip}\left(\frac{\pi_{\theta}(\mathbf{y} \mid \mathbf{x})}{\pi_{\text{old}}(\mathbf{y} \mid \mathbf{x})}, 1 - \epsilon, 1 + \epsilon\right)^{\text{detach}}$ .

242 Incorporating these improvements into Eq. 5, we arrive at the following FlowRL objective:

### FlowRL

$$247 \quad \mathcal{L}_{\text{FlowRL}} = w \cdot \left( \log Z_{\phi}(\mathbf{x}) + \frac{1}{|\mathbf{y}|} \log \pi_{\theta}(\mathbf{y} \mid \mathbf{x}) - \beta \hat{r}(\mathbf{x}, \mathbf{y}) - \frac{1}{|\mathbf{y}|} \log \pi_{\text{ref}}(\mathbf{y} \mid \mathbf{x}) \right)^2 \quad (6)$$

250 where the clipped importance weight  $w$  and normalized reward  $\hat{r}(\mathbf{x}, \mathbf{y})$  are defined as:

$$253 \quad w = \text{clip}\left(\frac{\pi_{\theta}(\mathbf{y} \mid \mathbf{x})}{\pi_{\text{old}}(\mathbf{y} \mid \mathbf{x})}, 1 - \epsilon, 1 + \epsilon\right)^{\text{detach}}, \quad \hat{r}_i = \frac{r_i - \text{mean}(\mathbf{r})}{\text{std}(\mathbf{r})}. \quad (7)$$

254 We use this objective to update the policy parameters  $\theta$  during training, and refer to this strategy as  
 255 *FlowRL*. Implementation details and theoretical analysis are provided in § 4 and § B, respectively.

## 4 EXPERIMENT SETTINGS

260 **Backbone Models.** There are two learnable modules in Eq. 6: the policy model  $\pi_{\theta}$  and the  
 261 partition function  $Z_{\phi}$ . For the policy model  $\pi_{\theta}$ , we use Qwen-2.5-7B/32B ([Team, 2024](#)) for  
 262 math tasks and DeepSeek-R1-Distill-Qwen-7B ([DeepSeek-AI, 2025](#)) for code tasks, re-  
 263 spectively. The reference model  $\pi_{\text{ref}}$  is the corresponding fixed pretrained model. For partition  
 264 function  $Z_{\phi}$ , following [Lee et al. \(2024\)](#), we use a randomly initialized 3-layer MLP with hidden  
 265 dimensions matching those of the base model. The input to  $Z_{\phi}$  is the mean of the language model's  
 266 hidden states after encoding the input  $\mathbf{x}$ , and the output is a scalar value. We detail the implemen-  
 267 tation of  $Z_{\phi}$  in § F. All training scripts are based on the veRL ([Sheng et al., 2024](#)). For the reward  
 268 function, following [Lee et al. \(2024\)](#), we set the hyperparameter  $\beta = 15$ .

269 <sup>2</sup>The substitution replaces  $\beta r(\mathbf{x}, \mathbf{y})$  in trajectory balance objective Eq. 3 with  $\beta r(\mathbf{x}, \mathbf{y}) + \log \pi_{\text{ref}}(\mathbf{y} \mid \mathbf{x})$   
 to incorporate the reference model constraint.

270 **Table 1: Results on math reasoning benchmarks.** We report Avg@16 accuracy with relative  
 271 improvements shown as subscripts. Positive gains are shown in **green** and negative changes in **red**.  
 272 FlowRL outperforms all baselines across both 7B and 32B model scales.  
 273

Models	AIME24	AIME25	AMC23	MATH500	Minerva	Olympiad	Avg
Qwen2.5-32B-Base, Max Response Len = 8K tokens							
Backbone	4.58	2.08	28.59	52.48	26.99	21.37	22.68
R++	14.79 <sub>+10.21</sub>	9.17 <sub>+7.08</sub>	52.65 <sub>+24.06</sub>	44.35 <sub>-8.13</sub>	17.37 <sub>-9.62</sub>	24.52 <sub>+3.15</sub>	27.14
PPO	26.87 <sub>+22.29</sub>	20.41 <sub>+18.33</sub>	76.40 <sub>+47.81</sub>	69.17 <sub>+16.69</sub>	28.79 <sub>+1.80</sub>	37.90 <sub>+16.53</sub>	43.25
GRPO	23.12 <sub>+18.54</sub>	14.58 <sub>+12.50</sub>	76.87 <sub>+48.28</sub>	61.60 <sub>+9.12</sub>	18.95 <sub>-8.04</sub>	34.94 <sub>+13.57</sub>	38.34
FlowRL	23.95 <sub>+19.37</sub>	21.87 <sub>+19.79</sub>	73.75 <sub>+45.16</sub>	80.75 <sub>+28.27</sub>	38.21 <sub>+11.22</sub>	51.83 <sub>+30.46</sub>	<b>48.39</b>
Qwen2.5-7B-Base, Max Response Len = 8K tokens							
Backbone	4.38	2.08	30.78	54.47	22.38	24.03	23.02
R++	11.04 <sub>+6.66</sub>	5.41 <sub>+3.33</sub>	66.71 <sub>+35.93</sub>	54.25 <sub>-0.22</sub>	24.37 <sub>+1.99</sub>	27.33 <sub>+3.30</sub>	31.52
PPO	9.38 <sub>+5.00</sub>	7.29 <sub>+5.21</sub>	63.43 <sub>+32.65</sub>	57.98 <sub>+3.51</sub>	26.53 <sub>+4.15</sub>	27.25 <sub>+3.22</sub>	31.98
GRPO	13.54 <sub>+9.16</sub>	9.79 <sub>+7.71</sub>	64.53 <sub>+33.75</sub>	57.05 <sub>+2.58</sub>	23.06 <sub>+0.68</sub>	26.88 <sub>+2.85</sub>	32.48
FlowRL	15.41 <sub>+11.03</sub>	10.83 <sub>+8.75</sub>	54.53 <sub>+23.75</sub>	66.96 <sub>+12.49</sub>	31.41 <sub>+9.03</sub>	34.61 <sub>+10.58</sub>	<b>35.63</b>

287 **Table 2: Results on code benchmarks.** We report metrics with relative improvements shown as  
 288 subscripts. Positive gains are shown in **green** and negative changes in **red**. FlowRL achieves the  
 289 strongest performance across all three benchmarks.  
 290

Models	LiveCodeBench		CodeForces		HumanEval+
	Avg@16	Pass@16	Rating	Percentile	
DeepSeek-R1-Distill-Qwen-7B, Max Response Len = 8K tokens					
Backbone	30.68	49.46	886.68	19.4%	80.90
R++	30.46 <sub>-0.22</sub>	52.68 <sub>+3.22</sub>	1208.03 <sub>+321.35</sub>	56.8% <sub>+37.4%</sub>	76.61 <sub>-4.29</sub>
PPO	35.10 <sub>+4.42</sub>	54.48 <sub>+5.02</sub>	1403.07 <sub>+516.39</sub>	73.7% <sub>+54.3%</sub>	82.32 <sub>+1.42</sub>
GRPO	32.75 <sub>+2.07</sub>	52.32 <sub>+2.86</sub>	1313.82 <sub>+427.14</sub>	67.1% <sub>+47.7%</sub>	80.13 <sub>-0.77</sub>
FlowRL	<b>37.43</b> <sub>+6.75</sub>	<b>56.27</b> <sub>+6.81</sub>	<b>1549.47</b> <sub>+662.79</sub>	<b>83.3%</b> <sub>+63.9%</sub>	<b>83.28</b> <sub>+2.38</sub>

300 **Baselines.** We compare our method against three representative reward-maximization RL base-  
 301 lines: REINFORCE++ (R++; Sutton et al., 1999b; Hu et al., 2025), PPO (Schulman et al., 2017),  
 302 and GRPO (Shao et al., 2024). All baselines follow the official veRL recipes, with consistent train-  
 303 ing configurations. For fair comparison, all methods use the same learning rate, batch size, and  
 304 training steps, and are evaluated at convergence using identical step counts.  
 305

306 **Training Configuration.** We experiment on both math and code domains. For the math domain,  
 307 we use the training set collected from DAPO (Yu et al., 2025b). For the code domain, we follow  
 308 the setup of DeepCoder (Luo et al., 2025), using their training set. For 7B model training, we  
 309 use a single node equipped with 8 NVIDIA H800 GPUs (80GB memory each). For 32B model  
 310 training, we scale to 4 nodes with 32 GPUs to accommodate the larger memory requirements. All  
 311 experiments use `max_prompt_length` = 2048 and `max_response_length` = 8192 across both  
 312 model sizes. We use a batch size of 512 for math reasoning tasks and 64 for code reasoning tasks.  
 313 We set the learning rate to 1e-6 and enable dynamic batch sizing in veRL for efficient training. For  
 314 GRPO and FlowRL, we configure `rollout_n` = 8, meaning each prompt generates 8 response  
 315 rollouts as the group size.  
 316

317 **Evaluation Configuration.** For the math domain, we evaluate on six challenging benchmarks:  
 318 AIME 2024/2025 (MAA, 2025), AMC 2023 (MAA, 2023), MATH-500 (Lightman et al., 2023a),  
 319 Minerva (Lewkowycz et al., 2022), and Olympiad (He et al., 2024). For the code domain, we eval-  
 320 uate on LiveCodeBench (Jain et al., 2024), CodeForces (Penedo et al., 2025), and HumanEval+ (Chen  
 321 et al., 2021). For all evaluation datasets, we perform 16 rollouts and report the average Pass@1  
 322 accuracy, denoted as Avg@16. We further report rating and percentile for Codeforces. During gen-  
 323 eration, we use sampling parameters of `temperature` = 0.6 and `top_p` = 0.95 for all evaluations.  
 The response length for evaluation is set to 8,192 tokens, consistent with the training configuration.

324  
325  
326  
327  
328  
329  
330  
331  
332  
333  
334  
335  
336  
337  
338  
339  
340  
341  
342  
343  
344  
345  
346  
347  
348  
349  
350  
351  
352  
353  
354  
355  
356  
357  
358  
359  
360  
361  
362  
363  
364  
365  
366  
367  
368  
369  
370  
371  
372  
373  
374  
375  
376  
377  
Table 3: Ablation study on FlowRL with Qwen2.5-7B as the base model. Avg@16 accuracy is reported across six math reasoning benchmarks. IS denotes importance sampling.

Method	AIME 2024	AIME 2025	AMC 2023	MATH-500	Minerva	Olympiad	Avg
FlowRL	15.41	10.83	54.53	66.96	31.41	34.61	35.63
w/o IS	6.25	7.91	41.40	56.97	22.19	25.52	26.71
Zhang et al. (2025a)	10.41	6.66	53.75	66.50	30.97	33.72	33.67

## 5 RESULTS

**Main Results.** Our experimental results, summarized in Table 1 and Table 2, demonstrate that FlowRL consistently outperforms all reward-maximization baselines across both math and code reasoning domains. Table 1 reports results on math reasoning benchmarks using both 7B and 32B base models, while Table 2 presents the corresponding results on code reasoning tasks. On math reasoning tasks, FlowRL achieves the highest average accuracy of 35.6% with the 7B model and 48.4% with the 32B model, surpassing PPO by 5.1% and GRPO by 10.1% on the 32B model. FlowRL shows strong improvements on challenging benchmarks like MATH-500 and Olympiad problems, demonstrating consistent gains across diverse mathematical domains. On code generation tasks, FlowRL achieves compelling improvements with the highest Avg@16 score of 37.43% on LiveCodeBench, a Codeforces rating of 1549.47 with 83.3% percentile ranking, and 83.28% accuracy on HumanEval+, outperforming all baselines across the board. These consistent performance gains across both domains and model scales provide strong empirical evidence that FlowRL’s flow-balanced optimization successfully enhances generalization. This improvement comes from promoting diverse solution exploration compared to previous reward-maximizing RL approaches.

**Ablation Studies.** We conduct ablation studies on importance sampling and the  $\beta$  hyperparameter. For importance sampling, we compared the performance with and without it, and implemented a combined loss approach proposed by Zhang et al. (2025a) that simultaneously optimizes both GFlowNets and PPO objectives. This combined loss focuses on optimizing diffusion models, and we adapt it to long CoT reasoning tasks for comparison. Table 3 demonstrates that importance sampling substantially improves FlowRL performance across all math reasoning benchmarks. Compared to Zhang et al. (2025a), using importance sampling as a trajectory-level ratio is more suitable than the combined loss of GFlowNets and PPO. The performance drop without importance sampling (from 35.63% to 26.71%) highlights the critical role of correcting for distribution mismatch between rollout generation and policy training. For the hyperparameter  $\beta$ , we conduct a series of parameter ablation studies, and Figure 3 shows that  $\beta = 15$  achieves optimal performance, with detailed results shown in Table 7.

## 6 ANALYSIS

**Diversity Analysis.** To assess solution diversity, we follow the approach of Yu et al. (2025a) and employ GPT-4o-mini (OpenAI, 2024) to evaluate all responses generated by each method on AIME 24/25. The evaluation prompt is shown in Appendix H. As shown in Figure 4, FlowRL achieves higher diversity scores compared to baseline methods. This demonstrates that FlowRL improves sample diversity compared to baselines, which tend to exhibit repetitive solution patterns. This diversity evaluation reveals significant differences in exploration patterns across methods. This nearly doubling of diversity score compared to the strongest baseline (PPO) indicates that FlowRL generates qualitatively different solution approaches rather than minor variations of the same strat-

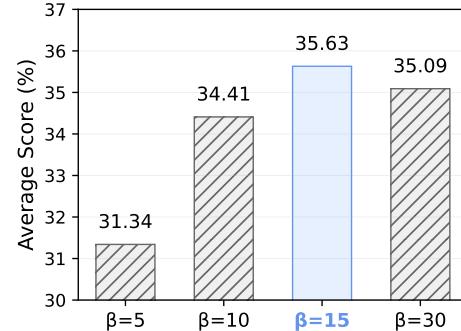


Figure 3: Ablation study on the  $\beta$  in FlowRL.  $\beta = 15$  (highlighted in blue) achieves the best performance.

378  
 379 Table 4: Case study comparing GRPO and FlowRL rollouts on an AIME problem. GRPO exhibits  
 380 repetitive patterns (AM-GM  $\times 3$ , identity loops  $\times 2$ ), while FlowRL follows a more diverse solution  
 381 path.

Content (boxed = actions; “ $\times k$ ” = repeated; “...” = omitted)				
Question	Let $\mathcal{B}$ be the set of rectangular boxes with surface area 54 and volume 23. Let $r$ be the radius of the smallest sphere that can contain each box in $\mathcal{B}$ . If $r^2 = \frac{p}{q}$ with $\gcd(p, q) = 1$ , find $p + q$ .			
GRPO	$\boxed{\text{“... denote } a, b, c \dots}}$ $\boxed{2(ab+bc+ca) = 54, abc = 23 \dots}$ $\boxed{d = \sqrt{a^2 + b^2 + c^2}, r = d/2 \dots}$ $\boxed{(a+b+c)^2 = a^2 + b^2 + c^2 + 2(ab+bc+ca) \dots}$ $\dots$ $\boxed{\text{AM-GM } \times 3: \text{ AM-GM (1) } \dots \text{ AM-GM (2) } \dots \text{ AM-GM (3) } \dots}$ $\boxed{(a+b+c)^3 \text{ identity loop } \times 2: \text{ loop (1) } \dots \text{ loop (2) } \dots \text{ } a = b = c \text{ (contradiction) } \dots}$ $\boxed{\text{back to } (a+b+c)^2 \dots \text{ no factorization } \dots}$			
FlowRL	$\boxed{\text{“... let } a, b, c \text{ with } 2(ab+bc+ca) = 54, abc = 23 \dots}$ $\boxed{d = \sqrt{a^2 + b^2 + c^2}, r = d/2 \dots}$ $\boxed{(a+b+c)^2 \Rightarrow a^2 + b^2 + c^2 = s^2 - 54 \dots}$ $\dots$ $\boxed{a = b \dots}$ $\boxed{a^3 - 27a + 46 = 0 \dots}$ $\dots$ $\boxed{\text{rational root } a = 2 \dots}$ $\dots$ $\boxed{\text{factor } (a - 2)(a^2 + 2a - 23) \dots}$ $\dots$ $\boxed{\text{branch } a = -1 + 2\sqrt{6} \dots}$ $\boxed{\text{back-sub } c = 23/a^2 \dots}$ $\boxed{a^2 + b^2 + c^2 = \frac{657}{16} \dots}$ $\boxed{r^2 = \frac{657}{64} \dots}$ $\boxed{\text{Answer 721} \dots}$			

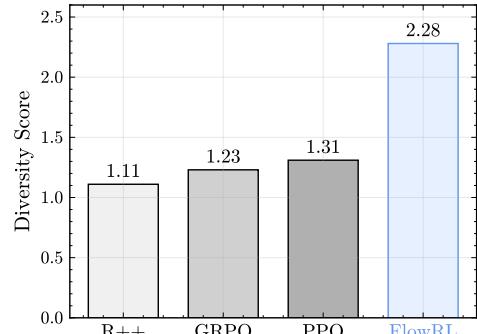
402 egy. The diversity analysis provides empirical validation of our core hypothesis that flow-balanced  
 403 optimization promotes mode coverage in complex reasoning tasks.

404  
 405 **Case Study.** Table 4 illustrates the behavioral differences between GRPO and FlowRL on a  
 406 representative AIME problem. GRPO exhibits repetitive  
 407 patterns, applying AM-GM three times and getting  
 408 stuck in identity loops, failing to solve the problem.  
 409 FlowRL explores more diverse actions: it sets  
 410  $a = b$ , derives a cubic equation, finds the rational  
 411 root, and reaches the correct answer. This shows  
 412 that FlowRL successfully avoids the repetitive ex-  
 413 ploration patterns. The contrast reveals fundamen-  
 414 tal differences in exploration strategies: GRPO’s  
 415 reward-maximizing approach leads to exploitation  
 416 of familiar techniques (AM-GM inequality) with-  
 417 out exploring alternatives, eventually reaching con-  
 418 tradictory conclusions like  $a = b = c$ . In con-  
 419 trast, FlowRL’s distribution-matching enables stra-  
 420 tegic decisions such as the symmetry assumption  $a =$   
 421  $b$ , which transforms the problem into a tractable cu-  
 422 bic equation  $a^3 - 27a + 46 = 0$ , allowing systematic solution through rational root testing and  
 423 polynomial factorization.

## 424 7 RELATED WORK

425 Our work relates to GFlowNets, Flow-Matching Policies, Length Normalization and KL Regular-  
 426 ization. We discuss three topics that relate most closely to our work in this section, and the other  
 427 topics are included in Appendix E.

428 429 **Reinforcement Learning for LLM Reasoning.** RL has emerged as a powerful approach for LLM  
 430 post-training on reasoning tasks (Sutton et al., 1999b; Schulman et al., 2017; Lightman et al., 2023b;



431 Figure 4: GPT-judged diversity scores on  
 432 rollouts of AIME 24/25 problems. FlowRL  
 433 generates more diverse solutions than R++,  
 434 GRPO, and PPO.

432 Shao et al., 2024; Guo et al., 2025). Most approaches employ reward-maximizing RL to optimize  
 433 expected cumulative returns. Entropy regularization (Haarnoja et al., 2018; Ahmed et al., 2019;  
 434 Cheng et al., 2025) is a classical technique for mitigating mode collapse by promoting diversity in  
 435 the policy’s output distribution, and has also been shown to enhance reasoning capabilities in var-  
 436 ious settings (Eysenbach & Levine, 2021; Chao et al., 2024). However, for long CoT reasoning,  
 437 the extended trajectory length (e.g., more than 8k tokens) makes it difficult for the regularization  
 438 signal to effectively influence reward-maximizing learning. Recent work (Cheng et al., 2025; Wang  
 439 et al., 2025; Cui et al., 2025; Dong et al., 2025) has discovered that training with more diverse or  
 440 high-entropy training data can further enhance training effectiveness. Compared to traditional en-  
 441 tropy regularization, the above methods explicitly increase the proportion of low-probability (i.e.,  
 442 high-entropy) tokens in the training data. In our work, we address the mode-collapse problem by  
 443 fundamentally shifting from reward maximization to reward distribution matching in our RL formu-  
 444 lation. See Appendix E for detailed comparisons.  
 445

446 **GFlowNets.** GFlowNets (Bengio et al., 2023a) represent a class of diversity-driven algorithms  
 447 designed to balance probability flows across states. They have rich connections to probabilistic  
 448 modeling methods (Zhang et al., 2022a;b; 2024a; Zimmermann et al., 2022; Malkin et al., 2023; Ma  
 449 et al.), and control methods (Pan et al., 2023b;c;d; Zhang et al., 2024b; Tiapkin et al., 2024). This  
 450 advantage has enabled GFlowNets to achieve successful applications in multiple downstream tasks,  
 451 such as molecular drug discovery (Jain et al., 2022; 2023b; Liu et al., 2022; Jain et al., 2023a; Shen  
 452 et al., 2023; Pan et al., 2023a; Kim et al., 2023; 2024), phylogenetic inference (Zhou et al., 2024),  
 453 and combinatorial optimization (Zhang et al., 2023a;b). For generative AI, GFlowNets provide a  
 454 powerful approach to align pretrained models in scenarios such as image generation (Zhang et al.,  
 455 2025a; Yun et al., 2025) and language model fine-tuning (Hu et al., 2024; Yu et al., 2025a; Lee  
 456 et al., 2024). Another line of work primarily focuses on the theoretical aspects of GFlowNets. Re-  
 457 cent theoretical studies have interpreted GFlowNets as solving a maximum entropy reinforcement  
 458 learning problem within a modified Markov Decision Process (MDP) (Tiapkin et al., 2024; Deleu  
 459 et al., 2024; Mohammadpour et al., 2024). These theoretical contributions have inspired us to en-  
 460 hance reinforcement learning from a more foundational standpoint using GFlowNets principles. A  
 461 comprehensive overview of GFlowNets theory can be found in Appendix C.  
 462

463 **Flow-Matching Policies.** Flow matching simplifies diffusion-based approaches by learning vector  
 464 fields that transport samples from prior to target distributions (Lipman et al., 2023). Recent work  
 465 has explored flow matching for policy optimization. McAllister et al. (2025) reformulates policy  
 466 optimization using advantage-weighted ratios from conditional flow matching loss, enabling flow-  
 467 based policy training without expensive likelihood computations. Pfrommer et al. (2025) explored  
 468 reward-weighted flow matching for improving policies beyond demonstration performance. Park  
 469 et al. (2025) uses a separate one-step policy to avoid unstable backpropagation through time when  
 470 training flow policies with RL. Zhang et al. (2025a) proposed a combined loss function integrating  
 471 PPO and GFlowNets to optimize diffusion model alignment. Lv et al. (2025) integrates flow-based  
 472 policy representation with Wasserstein regularized optimization for online reinforcement learning.  
 473 However, these approaches focus on continuous control, image generation, or vision-action models,  
 474 rather than addressing mode-collapse limitations in reward-maximizing RL. Inspired by flow match-  
 475 ing principles, our work improves upon RL training to enhance training stability while promoting  
 476 diverse solution exploration.  
 477

## 8 CONCLUSION

478 In this work, we introduce FlowRL, which transforms scalar rewards into normalized target distri-  
 479 butions using a learnable partition function and minimizes the reverse KL divergence between the  
 480 policy and target distribution. We demonstrate that this approach is theoretically equivalent to trajec-  
 481 tory balance objectives from GFlowNets and implicitly maximizes both reward and entropy, thereby  
 482 promoting diverse reasoning trajectories. To further address gradient explosion and sampling mis-  
 483 match issues in long CoT reasoning, we incorporate importance sampling and length normalization.  
 484 Through experiments on math and code reasoning benchmarks, FlowRL achieves consistent im-  
 485 provements across all tasks compared to GRPO and PPO. Our diversity analysis and case studies  
 486 confirm that FlowRL generates more varied solution approaches while avoiding repetitive patterns.  
 487

486 ETHICS STATEMENT  
487488 This work presents FlowRL, a reinforcement learning algorithm for improving reasoning in large  
489 language models. Our focus on mathematical and logical problem-solving directly supports ben-  
490 efitical applications in education, scientific research, and decision-support systems. We use estab-  
491 lished public benchmarks to ensure transparent and unbiased evaluation, and minimize computa-  
492 tional waste through efficient configurations, demonstrating our commitment to environmentally  
493 conscious and reproducible research.494  
495 REPRODUCIBILITY STATEMENT  
496497 We provide comprehensive details to ensure reproducibility: implementation specifics in Section 4  
498 (model architectures, training configurations, hyperparameters), complete algorithmic formulation  
499 in Eq. 6, experimental setup covering datasets and evaluation benchmarks, baseline implementa-  
500 tions following official verL recipes, and evaluation methodology. All mathematical formulations,  
501 implementation details, and experimental configurations necessary for reproduction are included in  
502 the paper.503  
504 REFERENCES  
505506 Zafarali Ahmed, Nicolas Le Roux, Mohammad Norouzi, and Dale Schuurmans. Understanding the  
507 impact of entropy on policy optimization. In *International conference on machine learning*, pp.  
508 151–160. PMLR, 2019.509 Mohammad Gheshlaghi Azar, Zhaoan Daniel Guo, Bilal Piot, Remi Munos, Mark Rowland,  
510 Michal Valko, and Daniele Calandriello. A general theoretical paradigm to understand learn-  
511 ing from human preferences. In *International Conference on Artificial Intelligence and Statistics*,  
512 pp. 4447–4455. PMLR, 2024.513 Brian R Bartoldson, Siddarth Venkatraman, James Diffenderfer, Moksh Jain, Tal Ben-Nun, Seanie  
514 Lee, Minsu Kim, Johan Obando-Ceron, Yoshua Bengio, and Bhavya Kailkhura. Trajectory bal-  
515 ance with asynchrony: Decoupling exploration and learning for fast, scalable ILM post-training.  
516 *arXiv preprint arXiv:2503.18929*, 2025.517 Emmanuel Bengio, Moksh Jain, Maksym Korablyov, Doina Precup, and Yoshua Bengio. Flow net-  
518 work based generative models for non-iterative diverse candidate generation. *Neural Information*  
519 *Processing Systems (NeurIPS)*, 2021.520 Yoshua Bengio, Salem Lahlou, Tristan Deleu, Edward J. Hu, Mo Tiwari, and Emmanuel Ben-  
521 gio. Gflownet foundations. *Journal of Machine Learning Research*, 24(210):1–55, 2023a. URL  
522 <http://jmlr.org/papers/v24/22-0364.html>.523 Yoshua Bengio, Salem Lahlou, Tristan Deleu, Edward J Hu, Mo Tiwari, and Emmanuel Bengio.  
524 Gflownet foundations. *The Journal of Machine Learning Research*, 24(1):10006–10060, 2023b.525 Chen-Hao Chao, Chien Feng, Wei-Fang Sun, Cheng-Kuang Lee, Simon See, and Chun-Yi Lee.  
526 Maximum entropy reinforcement learning via energy-based normalizing flow. *arXiv preprint*  
527 *arXiv:2405.13629*, 2024.528 Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde de Oliveira Pinto, Jared  
529 Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, Alex Ray, Raul Puri,  
530 Gretchen Krueger, Michael Petrov, Heidy Khlaaf, Girish Sastry, Pamela Mishkin, Brooke Chan,  
531 Scott Gray, Nick Ryder, Mikhail Pavlov, Alethea Power, Lukasz Kaiser, Mohammad Bavarian,  
532 Clemens Winter, Philippe Tillet, Felipe Petroski Such, Dave Cummings, Matthias Plappert,  
533 Fotios Chantzis, Elizabeth Barnes, Ariel Herbert-Voss, William Hebgen Guss, Alex Nichol, Alex  
534 Paino, Nikolas Tezak, Jie Tang, Igor Babuschkin, Suchir Balaji, Shantanu Jain, William Saunders,  
535 Christopher Hesse, Andrew N. Carr, Jan Leike, Josh Achiam, Vedant Misra, Evan Morikawa, Alec  
536 Radford, Matthew Knight, Miles Brundage, Mira Murati, Katie Mayer, Peter Welinder, Bob Mc-  
537 Grew, Dario Amodei, Sam McCandlish, Ilya Sutskever, and Wojciech Zaremba. Evaluating large  
538 language models trained on code, 2021.

540 Daixuan Cheng, Shaohan Huang, Xuekai Zhu, Bo Dai, Wayne Xin Zhao, Zhenliang Zhang, and  
 541 Furu Wei. Reasoning with exploration: An entropy perspective. [arXiv preprint arXiv:2506.14758](https://arxiv.org/abs/2506.14758),  
 542 2025.

543 Miruna Cretu, Charles Harris, Ilya Igashov, Arne Schneuing, Marwin Segler, Bruno Correia, Julien  
 544 Roy, Emmanuel Bengio, and Pietro Liò. Synflownet: Design of diverse and novel molecules with  
 545 synthesis constraints. [arXiv preprint arXiv:2405.01155](https://arxiv.org/abs/2405.01155), 2024.

546 Ganqu Cui, Yuchen Zhang, Jiacheng Chen, Lifan Yuan, Zhi Wang, Yuxin Zuo, Haozhan Li, Yuchen  
 547 Fan, Huayu Chen, Weize Chen, et al. The entropy mechanism of reinforcement learning for  
 548 reasoning language models. [arXiv preprint arXiv:2505.22617](https://arxiv.org/abs/2505.22617), 2025.

549 DeepSeek-AI. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning,  
 550 2025. URL <https://arxiv.org/abs/2501.12948>.

551 Tristan Deleu, Padideh Nouri, Nikolay Malkin, Doina Precup, and Yoshua Bengio. Discrete proba-  
 552 bility inference as control in multi-path environments. [arXiv preprint arXiv:2402.10309](https://arxiv.org/abs/2402.10309), 2024.

553 Guanting Dong, Hangyu Mao, Kai Ma, Licheng Bao, Yifei Chen, Zhongyuan Wang, Zhongxia  
 554 Chen, Jiazen Du, Huiyang Wang, Fuzheng Zhang, et al. Agentic reinforced policy optimization.  
 555 [arXiv preprint arXiv:2507.19849](https://arxiv.org/abs/2507.19849), 2025.

556 Yilun Du and Igor Mordatch. Implicit generation and modeling with energy based models. [Advances  
 557 in neural information processing systems](https://proceedings.neurips.cc/paper/2019/file/9f3956a2359a24415a5a222734855517.pdf), 32, 2019.

558 Benjamin Eysenbach and Sergey Levine. Maximum entropy rl (provably) solves some robust rl  
 559 problems. [arXiv preprint arXiv:2103.06257](https://arxiv.org/abs/2103.06257), 2021.

560 Leo Gao, John Schulman, and Jacob Hilton. Scaling laws for reward model overoptimization. In  
 561 [International Conference on Machine Learning](https://proceedings.mlr.press/v130/gao23a/gao23a.pdf), pp. 10835–10866. PMLR, 2023.

562 Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu,  
 563 Shirong Ma, Peiyi Wang, Xiao Bi, et al. Deepseek-r1: Incentivizing reasoning capability in llms  
 564 via reinforcement learning. [arXiv preprint arXiv:2501.12948](https://arxiv.org/abs/2501.12948), 2025.

565 Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. Soft actor-critic: Off-policy  
 566 maximum entropy deep reinforcement learning with a stochastic actor. In [International conference  
 567 on machine learning](https://proceedings.mlr.press/v80/haarnoja18a/haarnoja18a.pdf), pp. 1861–1870. Pmlr, 2018.

568 Chaoqun He, Renjie Luo, Yuzhuo Bai, Shengding Hu, Zhen Leng Thai, Junhao Shen, Jinyi Hu,  
 569 Xu Han, Yujie Huang, Yuxiang Zhang, et al. Olympiadbench: A challenging benchmark for  
 570 promoting agi with olympiad-level bilingual multimodal scientific problems. [arXiv preprint  
 571 arXiv:2402.14008](https://arxiv.org/abs/2402.14008), 2024.

572 Haoran He, Can Chang, Huazhe Xu, and Ling Pan. Looking backward: Retrospective backward syn-  
 573 thesis for goal-conditioned GFlownets. In [The Thirteenth International Conference on Learning  
 574 Representations](https://proceedings.mlr.press/v130/he23a/he23a.pdf), 2025. URL <https://openreview.net/forum?id=fNMKqyvuZT>.

575 Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and  
 576 Jacob Steinhardt. Measuring massive multitask language understanding. [arXiv preprint  
 577 arXiv:2009.03300](https://arxiv.org/abs/2009.03300), 2020.

578 Geoffrey E. Hinton, Peter Dayan, Brendan J. Frey, and R M Neal. The “wake-sleep” algorithm for  
 579 unsupervised neural networks. [Science](https://science.sciencemag.org/content/268/5214/1158), 268 5214:1158–61, 1995.

580 Edward J Hu, Moksh Jain, Eric Elmoznino, Younese Kaddar, Guillaume Lajoie, Yoshua Bengio,  
 581 and Nikolay Malkin. Amortizing intractable inference in large language models. [arXiv preprint  
 582 arXiv:2310.04363](https://arxiv.org/abs/2310.04363), 2023.

583 Edward J Hu, Moksh Jain, Eric Elmoznino, Younese Kaddar, Guillaume Lajoie, Yoshua Bengio,  
 584 and Nikolay Malkin. Amortizing intractable inference in large language models. In [The Twelfth  
 585 International Conference on Learning Representations](https://proceedings.mlr.press/v130/hu23a/hu23a.pdf), 2024. URL <https://openreview.net/forum?id=Ouj6p4ca60>.

594 Jian Hu, Jason Klein Liu, and Wei Shen. Reinforce++: An efficient rlhf algorithm with robustness  
 595 to both prompt and reward models, 2025. URL [https://arxiv.org/abs/2501.3262:32–33](https://arxiv.org/abs/2501.3262), 2025.  
 596

597 Moksh Jain, Emmanuel Bengio, Alex Hernandez-Garcia, Jarrid Rector-Brooks, Bonaventure F.P.  
 598 Dossou, Chanakya Ekbote, Jie Fu, Tianyu Zhang, Micheal Kilgour, Dinghuai Zhang, Lena  
 599 Simine, Payel Das, and Yoshua Bengio. Biological sequence design with GFlowNets.  
 600 International Conference on Machine Learning (ICML), 2022.

601 Moksh Jain, Tristan Deleu, Jason S. Hartford, Cheng-Hao Liu, Alex Hernández-García, and Yoshua  
 602 Bengio. Gflownets for ai-driven scientific discovery. ArXiv, abs/2302.00615, 2023a. URL  
 603 <https://api.semanticscholar.org/CorpusID:256459319>.  
 604

605 Moksh Jain, Sharath Chandra Raparthi, Alex Hernandez-Garcia, Jarrid Rector-Brooks, Yoshua  
 606 Bengio, Santiago Miret, and Emmanuel Bengio. Multi-objective GFlowNets. International  
 607 Conference on Machine Learning (ICML), 2023b.

608 Naman Jain, King Han, Alex Gu, Wen-Ding Li, Fanjia Yan, Tianjun Zhang, Sida Wang, Armando  
 609 Solar-Lezama, Koushik Sen, and Ion Stoica. Livecodebench: Holistic and contamination free  
 610 evaluation of large language models for code. arXiv preprint arXiv:2403.07974, 2024.  
 611

612 Koray Kavukcuoglu. Gemini 2.5: Our most intelligent AI model, 2025.  
 613 URL <https://blog.google/technology/google-deepmind/gemini-model-thinking-updates-march-2025/>. Google Blog (The Keyword),  
 614 Published Mar. 25, 2025.  
 615

616 Minsu Kim, Taeyoung Yun, Emmanuel Bengio, Dinghuai Zhang, Yoshua Bengio, Sungsoo Ahn,  
 617 and Jinkyoo Park. Local search gflownets. ArXiv, abs/2310.02710, 2023.  
 618

619 Minsu Kim, Joohwan Ko, Taeyoung Yun, Dinghuai Zhang, Ling Pan, Woochang Kim, Jinkyoo Park,  
 620 Emmanuel Bengio, and Yoshua Bengio. Learning to scale logits for temperature-conditional  
 621 gflownets, 2024.

622 Seanie Lee, Minsu Kim, Lynn Cherif, David Dobre, Juho Lee, Sung Ju Hwang, Kenji Kawaguchi,  
 623 Gauthier Gidel, Yoshua Bengio, Nikolay Malkin, et al. Learning diverse attacks on large language  
 624 models for robust red-teaming and safety tuning. arXiv preprint arXiv:2405.18540, 2024.  
 625

626 Aitor Lewkowycz, Anders Andreassen, David Dohan, Ethan Dyer, Henryk Michalewski,  
 627 Vinay Ramasesh, Ambrose Sloane, Cem Anil, Imanol Schlag, Theo Gutman-Solo,  
 628 Yuhuai Wu, Behnam Neyshabur, Guy Gur-Ari, and Vedant Misra. Solving quantitative  
 629 reasoning problems with language models. In S. Koyejo, S. Mohamed,  
 630 A. Agarwal, D. Belgrave, K. Cho, and A. Oh (eds.), Advances in Neural Information  
 631 Processing Systems, volume 35, pp. 3843–3857. Curran Associates, Inc., 2022. URL  
 632 [https://proceedings.neurips.cc/paper\\_files/paper/2022/file/18abbeef8cfe9203fdf9053c9c4fe191-Paper-Conference.pdf](https://proceedings.neurips.cc/paper_files/paper/2022/file/18abbeef8cfe9203fdf9053c9c4fe191-Paper-Conference.pdf).  
 633

634 Hunter Lightman, Vineet Kosaraju, Yura Burda, Harri Edwards, Bowen Baker, Teddy Lee, Jan  
 635 Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. Let’s verify step by step. arXiv preprint  
 636 arXiv:2305.20050, 2023a.  
 637

638 Hunter Lightman, Vineet Kosaraju, Yuri Burda, Harrison Edwards, Bowen Baker, Teddy Lee, Jan  
 639 Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. Let’s verify step by step. In The Twelfth  
 640 International Conference on Learning Representations, 2023b.

641 Yaron Lipman, Ricky T. Q. Chen, Heli Ben-Hamu, Maximilian Nickel, and Matthew Le. Flow  
 642 matching for generative modeling. In The Eleventh International Conference on Learning  
 643 Representations, 2023. URL <https://openreview.net/forum?id=PqvMRDCJT9t>.  
 644

645 Dianbo Liu, Moksh Jain, Bonaventure F. P. Dossou, Qianli Shen, Salem Lahlou, Anirudh Goyal,  
 646 Nikolay Malkin, Chris C. Emezue, Dinghuai Zhang, Nadhir Hassen, Xu Ji, Kenji Kawaguchi, and  
 647 Yoshua Bengio. Gflowout: Dropout with generative flow networks. In International Conference  
 648 on Machine Learning, 2022.

648 Mingjie Liu, Shizhe Diao, Jian Hu, Ximing Lu, Xin Dong, Hao Zhang, Alexander Bukharin,  
 649 Shaokun Zhang, Jiaqi Zeng, Makesh Narsimhan Sreedhar, et al. Scaling up rl: Unlocking di-  
 650 verse reasoning in llms via prolonged training. [arXiv preprint arXiv:2507.12507](https://arxiv.org/abs/2507.12507), 2025a.  
 651

652 Zhen Liu, Tim Z Xiao, , Weiyang Liu, Yoshua Bengio, and Dinghuai Zhang. Efficient diversity-  
 653 preserving diffusion alignment via gradient-informed gflownets. In [ICLR](https://iclr.cc), 2025b.  
 654

655 Zichen Liu, Changyu Chen, Wenjun Li, Penghui Qi, Tianyu Pang, Chao Du, Wee Sun Lee, and Min  
 656 Lin. Understanding r1-zero-like training: A critical perspective. [arXiv preprint arXiv:2503.20783](https://arxiv.org/abs/2503.20783),  
 657 2025c.  
 658

659 Michael Luo, Sijun Tan, Roy Huang, Xiaoxiang Shi, Rachel Xin, Colin Cai, Ameen Patel, Alpay  
 660 Ariyak, Qingyang Wu, Ce Zhang, Li Erran Li, Raluca Ada Popa, Ion Stoica, and Tianjun Zhang.  
 661 Deepcoder: A fully open-source 14b coder at o3-mini level, 2025. Notion Blog.  
 662

663 Lei Lv, Yunfei Li, Yu Luo, Fuchun Sun, Tao Kong, Jiafeng Xu, and Xiao Ma. Flow-based policy for  
 664 online reinforcement learning. [arXiv preprint arXiv:2506.12811](https://arxiv.org/abs/2506.12811), 2025.  
 665

666 Jiangyan Ma, Emmanuel Bengio, Yoshua Bengio, and Dinghuai Zhang. Baking symmetry into  
 667 gflownets.  
 668 MAA. American mathematics competitions - amc. <https://maa.org/>, 2023.  
 669

670 MAA. American invitational mathematics examination - aime. <https://maa.org/>, 2025.  
 671

672 Kanika Madan, Jarrid Rector-Brooks, Maksym Korablyov, Emmanuel Bengio, Moksh Jain, An-  
 673 drei Cristian Nica, Tom Bosc, Yoshua Bengio, and Nikolay Malkin. Learning gflownets from  
 674 partial episodes for improved convergence and stability. In [International Conference on Machine](https://icml.cc)  
 675 [Learning](https://icml.cc), pp. 23467–23483. PMLR, 2023.  
 676

677 Nikolay Malkin, Moksh Jain, Emmanuel Bengio, Chen Sun, and Yoshua Bengio. Trajectory balance:  
 678 Improved credit assignment in gflownets. [Advances in Neural Information Processing Systems](https://advancesinneurips.org),  
 679 35:5955–5967, 2022.  
 680

681 Nikolay Malkin, Salem Lahlou, Tristan Deleu, Xu Ji, Edward Hu, Katie Everett, Dinghuai Zhang,  
 682 and Yoshua Bengio. GFlowNets and variational inference. [International Conference on Learning](https://iclr.cc)  
 683 [Representations \(ICLR\)](https://iclr.cc), 2023.  
 684

685 David McAllister, Songwei Ge, Brent Yi, Chung Min Kim, Ethan Weber, Hongsuk Choi, Haiwen  
 686 Feng, and Angjoo Kanazawa. Flow matching policy gradients. [arXiv preprint arXiv:2507.21053](https://arxiv.org/abs/2507.21053),  
 687 2025.  
 688

689 Sobhan Mohammadpour, Emmanuel Bengio, Emma Freijinger, and Pierre-Luc Bacon. Maximum  
 690 entropy gflownets with soft q-learning. In [International Conference on Artificial Intelligence and](https://icml.cc)  
 691 [Statistics](https://icml.cc), pp. 2593–2601. PMLR, 2024.  
 692

693 OpenAI. Gpt-4o mini. <https://openai.com/index/>, 2024. Accessed: 2024.  
 694

695 Alexander Pan, Kush Bhatia, and Jacob Steinhardt. The effects of reward misspecification: Mapping  
 696 and mitigating misaligned models. [arXiv preprint arXiv:2201.03544](https://arxiv.org/abs/2201.03544), 2022.  
 697

698 Ling Pan, Moksh Jain, Kanika Madan, and Yoshua Bengio. Pre-training and fine-tuning generative  
 699 flow networks, 2023a.  
 700

701 Ling Pan, Nikolay Malkin, Dinghuai Zhang, and Yoshua Bengio. Better training of GFlowNets with  
 702 local credit and incomplete trajectories. [International Conference on Machine Learning \(ICML\)](https://icml.cc),  
 703 2023b.  
 704

705 Ling Pan, Dinghuai Zhang, Aaron Courville, Longbo Huang, and Yoshua Bengio. Generative aug-  
 706 mented flow networks. [International Conference on Learning Representations \(ICLR\)](https://iclr.cc), 2023c.  
 707

708 Ling Pan, Dinghuai Zhang, Moksh Jain, Longbo Huang, and Yoshua Bengio. Stochastic generative  
 709 flow networks. [Uncertainty in Artificial Intelligence \(UAI\)](https://uai.org), 2023d.  
 710

702 Seohong Park, Qiyang Li, and Sergey Levine. Flow q-learning. In *Forty-second International*  
 703 *Conference on Machine Learning*, 2025. URL <https://openreview.net/forum?id=KVf2SFL1pi>.

704

705 Guilherme Penedo, Anton Lozhkov, Hynek Kydlíček, Loubna Ben Allal, Edward Beeching,  
 706 Agustín Piqueres Lajarín, Quentin Gallouédec, Nathan Habib, Lewis Tunstall, and Leandro von  
 707 Werra. Codeforces. <https://huggingface.co/datasets/open-rl/codeforces>,  
 708 2025.

709

710 Samuel Pfrommer, Yixiao Huang, and Somayeh Sojoudi. Reinforcement learning for flow-matching  
 711 policies. [arXiv preprint arXiv:2507.15073](https://arxiv.org/abs/2507.15073), 2025.

712

713 Abhinav Rastogi, Albert Q Jiang, Andy Lo, Gabrielle Berrada, Guillaume Lample, Jason Rute, Joep  
 714 Barmentlo, Karmesh Yadav, Kartik Khandelwal, Khyathi Raghavi Chandu, et al. Magistral. [arXiv](https://arxiv.org/abs/2506.10910)  
 715 preprint arXiv:2506.10910, 2025.

716

717 David Rein, Betty Li Hou, Asa Cooper Stickland, Jackson Petty, Richard Yuanzhe Pang, Julien Di-  
 718 rani, Julian Michael, and Samuel R Bowman. Gpqa: A graduate-level google-proof q&a bench-  
 719 mark. In *First Conference on Language Modeling*, 2024.

720

721 John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy  
 722 optimization algorithms. [arXiv preprint arXiv:1707.06347](https://arxiv.org/abs/1707.06347), 2017.

723

724 Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang,  
 725 Mingchuan Zhang, YK Li, Y Wu, et al. Deepseekmath: Pushing the limits of mathematical  
 726 reasoning in open language models. [arXiv preprint arXiv:2402.03300](https://arxiv.org/abs/2402.03300), 2024.

727

728 Max W. Shen, Emmanuel Bengio, Ehsan Hajiramezanali, Andreas Loukas, Kyunghyun Cho,  
 729 and Tommaso Biancalani. Towards understanding and improving gflownet training. [ArXiv](https://arxiv.org/abs/2305.07170),  
 730 abs/2305.07170, 2023. URL <https://api.semanticscholar.org/CorpusID:258676487>.

731

732 Guangming Sheng, Chi Zhang, Zilingfeng Ye, Xibin Wu, Wang Zhang, Ru Zhang, Yanghua Peng,  
 733 Haibin Lin, and Chuan Wu. Hybridflow: A flexible and efficient rlhf framework. [arXiv preprint](https://arxiv.org/abs/2409.19256)  
 734 [arXiv: 2409.19256](https://arxiv.org/abs/2409.19256), 2024.

735

736 Joar Skalse, Nikolaus Howe, Dmitrii Krasheninnikov, and David Krueger. Defining and charac-  
 737 terizing reward gaming. *Advances in Neural Information Processing Systems*, 35:9460–9471,  
 738 2022.

739

740 Richard S Sutton, Andrew G Barto, et al. Reinforcement learning. *Journal of Cognitive*  
 741 *Neuroscience*, 11(1):126–134, 1999a.

742

743 Richard S Sutton, David McAllester, Satinder Singh, and Yishay Mansour. Policy gradient  
 744 methods for reinforcement learning with function approximation. In S. Solla, T. Leen, and  
 745 K. Müller (eds.), *Advances in Neural Information Processing Systems*, volume 12. MIT Press,  
 746 1999b. URL [https://proceedings.neurips.cc/paper\\_files/paper/1999/file/464d828b85b0bed98e80ade0a5c43b0f-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/1999/file/464d828b85b0bed98e80ade0a5c43b0f-Paper.pdf).

747

748 Kimi Team, Angang Du, Bofei Gao, Bowei Xing, Changjiu Jiang, Cheng Chen, Cheng Li, Chenjun  
 749 Xiao, Chenzhuang Du, Chonghua Liao, et al. Kimi k1. 5: Scaling reinforcement learning with  
 750 llms. [arXiv preprint arXiv:2501.12599](https://arxiv.org/abs/2501.12599), 2025.

751

752 Qwen Team. Qwen2.5: A party of foundation models, September 2024. URL <https://qwenlm.github.io/blog/qwen2.5/>.

753

754 Daniil Tiapkin, Nikita Morozov, Alexey Naumov, and Dmitry P Vetrov. Generative flow networks  
 755 as entropy-regularized rl. In *International Conference on Artificial Intelligence and Statistics*, pp.  
 4213–4221. PMLR, 2024.

756

757 Shenzhi Wang, Le Yu, Chang Gao, Chujie Zheng, Shixuan Liu, Rui Lu, Kai Dang, Xionghui Chen,  
 758 Jianxin Yang, Zhenru Zhang, et al. Beyond the 80/20 rule: High-entropy minority tokens drive  
 759 effective reinforcement learning for llm reasoning. [arXiv preprint arXiv:2506.01939](https://arxiv.org/abs/2506.01939), 2025.

756 Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny  
 757 Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. *Advances in*  
 758 *neural information processing systems*, 35:24824–24837, 2022.

759

760 Fangxu Yu, Lai Jiang, Haoqiang Kang, Shibo Hao, and Lianhui Qin. Flow of reasoning: Training  
 761 llms for divergent reasoning with minimal examples. In *Forty-second International Conference*  
 762 *on Machine Learning*, 2025a.

763 Qiyi Yu, Zheng Zhang, Ruofei Zhu, Yufeng Yuan, Xiaochen Zuo, Yu Yue, Tiantian Fan, Gaohong  
 764 Liu, Lingjun Liu, Xin Liu, et al. Dapo: An open-source llm reinforcement learning system at  
 765 scale. *arXiv preprint arXiv:2503.14476*, 2025b.

766 Taeyoung Yun, Dinghuai Zhang, Jinkyoo Park, and Ling Pan. Learning to sample effective and  
 767 diverse prompts for text-to-image generation. In *Proceedings of the Computer Vision and Pattern*  
 768 *Recognition Conference*, pp. 23625–23635, 2025.

769

770 Eric Zelikman, Yuhuai Wu, Jesse Mu, and Noah Goodman. Star: Bootstrapping reasoning with  
 771 reasoning. *Advances in Neural Information Processing Systems*, 35:15476–15488, 2022.

772 David W. Zhang, Corrado Rainone, Markus F. Peschl, and Roberto Bondesan. Robust scheduling  
 773 with gflownets. *ArXiv*, abs/2302.05446, 2023a. URL <https://api.semanticscholar.org/CorpusID:256827133>.

774

775 Dinghuai Zhang, Ricky T. Q. Chen, Nikolay Malkin, and Yoshua Bengio. Unifying generative  
 776 models with GFlowNets and beyond. *arXiv preprint arXiv:2209.02606v2*, 2022a.

777

778 Dinghuai Zhang, Nikolay Malkin, Zhen Liu, Alexandra Volokhova, Aaron Courville, and Yoshua  
 779 Bengio. Generative flow networks for discrete probabilistic modeling. *International Conference*  
 780 *on Machine Learning (ICML)*, 2022b.

781

782 Dinghuai Zhang, Hanjun Dai, Nikolay Malkin, Aaron C. Courville, Yoshua Bengio, and Ling Pan.  
 783 Let the flows tell: Solving graph combinatorial optimization problems with gflownets. *ArXiv*,  
 784 abs/2305.17010, 2023b.

785

786 Dinghuai Zhang, Ricky T. Q. Chen, Cheng-Hao Liu, Aaron Courville, and Yoshua Bengio. Diffu-  
 787 sion generative flow samplers: Improving learning signals through partial trajectory optimization,  
 788 2024a.

789

790 Dinghuai Zhang, Ling Pan, Ricky T. Q. Chen, Aaron Courville, and Yoshua Bengio. Distributional  
 791 gflownets with quantile flows, 2024b.

792

793 Dinghuai Zhang, Yizhe Zhang, Jiatao Gu, Ruixiang ZHANG, Joshua M. Susskind, Navdeep Jaitly,  
 794 and Shuangfei Zhai. Improving GFlownets for text-to-image diffusion alignment. *Transactions on*  
 795 *Machine Learning Research*, 2025a. ISSN 2835-8856. URL <https://openreview.net/forum?id=XDbY3qhM42>.

796

797 Kaiyan Zhang, Yuxin Zuo, Bingxiang He, Youbang Sun, Runze Liu, Che Jiang, Yuchen Fan, Kai  
 798 Tian, Guoli Jia, Pengfei Li, et al. A survey of reinforcement learning for large reasoning models.  
 799 *arXiv preprint arXiv:2509.08827*, 2025b.

800

801 Xiaojiang Zhang, Jinghui Wang, Zifei Cheng, Wenhao Zhuang, Zheng Lin, Minglei Zhang, Shaojie  
 802 Wang, Yinghan Cui, Chao Wang, Junyi Peng, et al. Srp0: A cross-domain implementation of  
 803 large-scale reinforcement learning on llm. *arXiv preprint arXiv:2504.14286*, 2025c.

804

805 Chujie Zheng, Shixuan Liu, Mingze Li, Xiong-Hui Chen, Bowen Yu, Chang Gao, Kai Dang,  
 806 Yuqiong Liu, Rui Men, An Yang, et al. Group sequence policy optimization. *arXiv preprint*  
 807 *arXiv:2507.18071*, 2025.

808

809 Mingyang Zhou, Zichao Yan, Elliot Layne, Nikolay Malkin, Dinghuai Zhang, Moksh Jain, Mathieu  
 810 Blanchette, and Yoshua Bengio. Phylogfn: Phylogenetic inference with generative flow networks,  
 811 2024.

812

813 Heiko Zimmermann, Fredrik Lindsten, J.-W. van de Meent, and Christian Andersson Naesseth.  
 814 A variational perspective on generative flow networks. *ArXiv*, abs/2210.07992, 2022. URL  
 815 <https://api.semanticscholar.org/CorpusID:252907672>.

810 A PROOF OF PROPOSITION 1  
811812 We begin by analyzing the gradient of the Kullback–Leibler (KL) divergence between the policy  
813  $\pi_\theta(\mathbf{y} \mid \mathbf{x})$  and the target reward distribution  $\frac{\exp(\beta r(\mathbf{x}, \mathbf{y}))}{Z_\phi(\mathbf{x})}$ .  
814

815 
$$\begin{aligned} & \nabla_\theta D_{\text{KL}} \left( \pi_\theta(\mathbf{y} \mid \mathbf{x}) \parallel \frac{\exp(\beta r(\mathbf{x}, \mathbf{y}))}{Z_\phi(\mathbf{x})} \right) \\ &= \nabla_\theta \int \pi_\theta(\mathbf{y} \mid \mathbf{x}) \log \left[ \frac{\pi_\theta(\mathbf{y} \mid \mathbf{x}) \cdot Z_\phi(\mathbf{x})}{\exp(\beta r(\mathbf{x}, \mathbf{y}))} \right] d\mathbf{y} \\ &= \int \nabla_\theta \pi_\theta(\mathbf{y} \mid \mathbf{x}) \log \left[ \frac{Z_\phi(\mathbf{x}) \pi_\theta(\mathbf{y} \mid \mathbf{x})}{\exp(\beta r(\mathbf{x}, \mathbf{y}))} \right] d\mathbf{y} + \int \pi_\theta(\mathbf{y} \mid \mathbf{x}) \nabla_\theta \log \left[ \frac{Z_\phi(\mathbf{x}) \pi_\theta(\mathbf{y} \mid \mathbf{x})}{\exp(\beta r(\mathbf{x}, \mathbf{y}))} \right] d\mathbf{y} \\ &= \int \pi_\theta(\mathbf{y} \mid \mathbf{x}) \nabla_\theta \log \pi_\theta(\mathbf{y} \mid \mathbf{x}) \log \left[ \frac{Z_\phi(\mathbf{x}) \pi_\theta(\mathbf{y} \mid \mathbf{x})}{\exp(\beta r(\mathbf{x}, \mathbf{y}))} \right] d\mathbf{y} + \underbrace{\int \pi_\theta(\mathbf{y} \mid \mathbf{x}) \nabla_\theta \log \pi_\theta(\mathbf{y} \mid \mathbf{x}) d\mathbf{y}}_{= \nabla_\theta \int \pi_\theta(\mathbf{y} \mid \mathbf{x}) d\mathbf{y} = \nabla_\theta 1 = 0} \\ &= \int \pi_\theta(\mathbf{y} \mid \mathbf{x}) \nabla_\theta \log \pi_\theta(\mathbf{y} \mid \mathbf{x}) \log \left[ \frac{Z_\phi(\mathbf{x}) \pi_\theta(\mathbf{y} \mid \mathbf{x})}{\exp(\beta r(\mathbf{x}, \mathbf{y}))} \right] d\mathbf{y} \\ &= \mathbb{E}_{\mathbf{y} \sim \pi_\theta(\cdot \mid \mathbf{x})} \left[ \log \left( \frac{Z_\phi(\mathbf{x}) \pi_\theta(\mathbf{y} \mid \mathbf{x})}{\exp(\beta r(\mathbf{x}, \mathbf{y}))} \right) \cdot \nabla_\theta \log \pi_\theta(\mathbf{y} \mid \mathbf{x}) \right] \end{aligned} \quad (8)$$

830 Next, consider the trajectory balance objective used in GFlowNets learning (Bengio et al., 2023b;  
831 Lee et al., 2024; Bartoldson et al., 2025), defined as:

832 
$$\mathcal{L}(\mathbf{y}, \mathbf{x}; \theta) = \left( \log \frac{Z_\phi(\mathbf{x}) \pi_\theta(\mathbf{y} \mid \mathbf{x})}{\exp(\beta r(\mathbf{x}, \mathbf{y}))} \right)^2. \quad (9)$$

833 Taking the gradient of this objective with respect to  $\theta$  yields:  
834

835 
$$\nabla_\theta \mathcal{L}(\theta) = 2 \cdot \mathbb{E}_{\mathbf{y} \sim \pi_\theta(\cdot \mid \mathbf{x})} \left[ \left( \log \frac{Z_\phi(\mathbf{x}) \cdot \pi_\theta(\mathbf{y} \mid \mathbf{x})}{\exp(\beta r(\mathbf{x}, \mathbf{y}))} \right) \cdot \nabla_\theta \log \pi_\theta(\mathbf{y} \mid \mathbf{x}) \right] \quad (10)$$

836 Thus, minimizing the KL divergence is equivalent (up to a constant) to minimizing the trajectory  
837 balance loss, confirming Proposition 1.838 B THEORETICAL ANALYSIS  
839840 We conduct an interpretation of FlowRL that clarifies the role of each component in the objective.  
841842 **Proposition 5.** *Minimizing the KL divergence in Eq. 5 is equivalent (in terms of gradients) to jointly  
843 maximizing reward and policy entropy:*

844 
$$\max_\theta \mathbb{E}_{\mathbf{y} \sim \pi_\theta} \left[ \underbrace{\beta r(\mathbf{x}, \mathbf{y})}_{\text{reward}} - \log Z_\phi(\mathbf{x}) + \log \pi_{\text{ref}}(\mathbf{y} \mid \mathbf{x}) \right] + \underbrace{\mathcal{H}(\pi_\theta)}_{\text{entropy}}. \quad (11)$$

845 **Remark 6 (FlowRL beyond reward maximization).** Proposition 5 reveals that FlowRL can be inter-  
846 preted as jointly maximizing expected reward and policy entropy. This formulation encourages the  
847 policy to explore a broader set of high-quality solutions, enabling more diverse and generalizable  
848 behaviors on reasoning tasks. Our interpretation also aligns with prior work that views GFlowNets  
849 training as a form of maximum entropy RL (Mohammadpour et al., 2024; Deleu et al., 2024).  
850851 The proof of Proposition 5 is provided as below.  
852853 Recall from Eq. 3 and Eq. 5 that the FlowRL objective is sourced from the minimization of a KL  
854 divergence:  
855

856 
$$D_{\text{KL}} \left( \pi_\theta(\mathbf{y} \mid \mathbf{x}) \parallel \frac{\exp(\beta r(\mathbf{x}, \mathbf{y})) \cdot \pi_{\text{ref}}(\mathbf{y} \mid \mathbf{x})}{Z_\phi(\mathbf{x})} \right) = \int \pi_\theta(\mathbf{y} \mid \mathbf{x}) \log \left[ \frac{Z_\phi(\mathbf{x}) \pi_\theta(\mathbf{y} \mid \mathbf{x})}{\exp(\beta r(\mathbf{x}, \mathbf{y})) \cdot \pi_{\text{ref}}(\mathbf{y} \mid \mathbf{x})} \right] d\mathbf{y} \quad (12)$$

864 Rearranging the terms, we obtain:  
 865

$$\begin{aligned}
 866 \quad & \arg \min_{\theta} D_{\text{KL}}\left(\pi_{\theta}(\mathbf{y} \mid \mathbf{x}) \parallel \frac{\exp (\beta r(\mathbf{x}, \mathbf{y})) \cdot \pi_{\text{ref}}(\mathbf{y} \mid \mathbf{x})}{Z_{\phi}(\mathbf{x})}\right) \\
 867 \quad & = \arg \min_{\theta} \int \pi_{\theta}(\mathbf{y} \mid \mathbf{x}) \log \left[\frac{Z_{\phi}(\mathbf{x}) \pi_{\theta}(\mathbf{y} \mid \mathbf{x})}{\exp (\beta r(\mathbf{x}, \mathbf{y})) \cdot \pi_{\text{ref}}(\mathbf{y} \mid \mathbf{x})}\right] d \mathbf{y} \\
 868 \quad & = \arg \max_{\theta}\left\{\mathbb{E}_{\mathbf{y} \sim \pi_{\theta}(\cdot \mid \mathbf{x})} \log \left[\frac{\exp (\beta r(\mathbf{x}, \mathbf{y})) \cdot \pi_{\text{ref}}(\mathbf{y} \mid \mathbf{x})}{Z_{\phi}(\mathbf{x})}\right]-\int \pi_{\theta}(\mathbf{y} \mid \mathbf{x}) \log \pi_{\theta}(\mathbf{y} \mid \mathbf{x}) d \mathbf{y}\right\} \\
 869 \quad & = \arg \max_{\theta}\left\{\mathbb{E}_{\mathbf{y} \sim \pi_{\theta}(\cdot \mid \mathbf{x})} \log \left[\frac{\exp (\beta r(\mathbf{x}, \mathbf{y})) \cdot \pi_{\text{ref}}(\mathbf{y} \mid \mathbf{x})}{Z_{\phi}(\mathbf{x})}\right]+\mathcal{H}(\pi_{\theta})\right\}
 \end{aligned} \tag{13}$$

870  
 871 Finally, we express the FlowRL objective in its compact form:  
 872

$$\max_{\theta} \mathbb{E}_{\mathbf{y} \sim \pi_{\theta}(\cdot \mid \mathbf{x})}\left[\underbrace{\beta r(\mathbf{x}, \mathbf{y})}_{\text {reward }}-\underbrace{\log Z_{\phi}(\mathbf{x})}_{\text {normalization }}+\underbrace{\log \pi_{\text {ref }}(\mathbf{y} \mid \mathbf{x})}_{\text {reference model constraint }}\right]+\underbrace{\mathcal{H}(\pi_{\theta})}_{\text {entropy }} . \tag{14}$$

873  
 874 Therefore, minimizing the FlowRL objective can be interpreted as jointly maximizing reward and  
 875 entropy, while also aligning the policy with a structured prior. The reward term drives task performance,  
 876 while the normalization term  $Z_{\phi}(\mathbf{x})$  ensures consistency with a properly normalized target  
 877 distribution. This encourages the policy  $\pi_{\theta}$  to cover the entire reward-weighted distribution rather  
 878 than collapsing to a few high-reward modes. The reference policy  $\pi_{\text{ref}}$  provides inductive bias that  
 879 regularizes the policy toward desirable structures, and the entropy term  $\mathcal{H}(\pi_{\theta})$  encourages diversity  
 880 in sampled solutions. Together, these components promote better generalization of FlowRL.  
 881

## C GFLOWNETS

882  
 883 We follow the notation of (Madan et al., 2023; He et al., 2025) to introduce the fundamentals of  
 884 GFNets. Let  $\mathcal{X}$  denote the compositional objects and  $R$  be a reward function that assigns non-  
 885 negative values to each object  $x \in \mathcal{X}$ . GFNets aim to learn a sequential, constructive sampling  
 886 policy  $\pi$  that generates objects  $x$  with probabilities proportional to their rewards, i.e.,  $\pi(x) \propto R(x)$ .  
 887 This process can be represented as a directed acyclic graph (DAG)  $\mathcal{G}=(\mathcal{S}, \mathcal{A})$ , where the vertices  
 888  $s \in \mathcal{S}$  are referred to as *states*, and the directed edges  $(u \rightarrow v) \in \mathcal{A}$  are called *actions*. The  
 889 generation of an object  $x \in \mathcal{X}$  corresponds to a complete trajectory  $\tau=(s_0 \rightarrow \cdots \rightarrow s_n) \in \mathcal{T}$   
 890 within the DAG, beginning at the initial state  $s_0$  and ending at a terminal state  $s_n \in \mathcal{X}$ . The state flow  
 891  $F(s)$  is defined as a non-negative weight assigned to each state  $s \in \mathcal{S}$ . The forward policy  $P_F(s' \mid s)$   
 892 specifies the transition probability to a child state  $s'$ , while the backward policy  $P_B(s \mid s')$  specifies  
 893 the transition probability to a parent state  $s$ . To this end, detailed balance objective enforces local  
 894 flow consistency across every edge  $(s \rightarrow s') \in \mathcal{A}$ :  
 895

$$\forall(s \rightarrow s') \in \mathcal{A}, \quad F_{\theta}(s) P_F(s' \mid s ; \theta)=F_{\theta}(s') P_B(s \mid s' ; \theta) . \tag{15}$$

896 To achieve this flow consistency, GFNets employ training objectives at different levels of granu-  
 897 larity, including detailed balance (Bengio et al., 2023b), trajectory balance (Malkin et al., 2022), and  
 898 sub-trajectory balance (Madan et al., 2023). Leveraging their diversity-seeking behavior, GFNets  
 899 have been successfully applied across a range of domains, including molecule generation (Cretu  
 900 et al., 2024), diffusion fine-tuning (Liu et al., 2025b; Zhang et al., 2025a), and amortized reasoning  
 901 (Hu et al., 2024; Yu et al., 2025a). Among various training objective in GFNets, trajectory  
 902 balance maintains flow consistency at the trajectory level, defined as:  
 903

$$Z_{\theta} \prod_{t=1}^n P_F\left(s_t \mid s_{t-1} ; \theta\right)=R(x) \prod_{t=1}^n P_B\left(s_{t-1} \mid s_t ; \theta\right) . \tag{16}$$

904  
 905 Furthermore, sub-trajectory balance achieves local balance on arbitrary subpaths  $\tau_{i:j}=\{s_i \rightarrow$   
 906  $\cdots \rightarrow s_j\}$ , offering a more stable and less biased learning signal. We build on trajectory balance  
 907 to extend our KL-based objective through a gradient-equivalence formulation (Prop. 1), and further  
 908 improve it to better support long CoT reasoning in RL.  
 909

918  
919 Table 5: Math reasoning performance (Avg@64) at temperature = 0.6. Relative improvements are  
920 shown as subscripts, with positive gains in **green** and negative changes in **red**. FlowRL consistently  
921 outperforms all baselines and achieves the best average score under this low-temperature setting.  
922

Models	AIME 2024	AIME 2025	AMC 2023	MATH-500	Minerva	Olympiad	Avg
Qwen2.5-7B Base Model							
Backbone	4.37	2.08	30.78	54.48	22.38	24.02	23.02
R++	10.57 <sub>+6.20</sub>	5.10 <sub>+3.02</sub>	66.02 <sub>+35.24</sub>	54.29 <sub>-0.19</sub>	24.47 <sub>+2.09</sub>	27.30 <sub>+3.28</sub>	31.29
PPO	9.95 <sub>+5.58</sub>	7.34 <sub>+5.26</sub>	63.63 <sub>+32.85</sub>	57.72 <sub>+3.24</sub>	26.22 <sub>+3.84</sub>	27.35 <sub>+3.33</sub>	32.03
GRPO	14.01 <sub>+9.64</sub>	10.73 <sub>+8.65</sub>	64.10 <sub>+33.32</sub>	57.41 <sub>+2.93</sub>	23.17 <sub>+0.79</sub>	27.11 <sub>+3.09</sub>	32.76
FlowRL	14.32 <sub>+9.95</sub>	10.05 <sub>+7.97</sub>	55.08 <sub>+24.30</sub>	66.78 <sub>+12.30</sub>	31.52 <sub>+9.14</sub>	34.60 <sub>+10.58</sub>	<b>35.39</b>

930  
931 Table 6: Math reasoning performance (Avg@64) at temperature = 1.0. Relative improvements  
932 are shown as subscripts, with positive gains in **green**. FlowRL maintains robust performance under  
933 higher generation randomness and continues to outperform all baselines on average.  
934

Models	AIME 2024	AIME 2025	AMC 2023	MATH-500	Minerva	Olympiad	Avg
Qwen2.5-7B Base Model							
Backbone	3.39	1.51	23.90	45.18	16.98	18.27	18.20
R++	10.63 <sub>+7.24</sub>	4.63 <sub>+3.12</sub>	66.99 <sub>+43.09</sub>	54.36 <sub>+9.18</sub>	23.89 <sub>+6.91</sub>	26.65 <sub>+8.38</sub>	31.19
PPO	10.52 <sub>+7.13</sub>	6.51 <sub>+5.00</sub>	63.04 <sub>+39.14</sub>	57.46 <sub>+12.28</sub>	25.91 <sub>+8.93</sub>	27.16 <sub>+8.89</sub>	31.77
GRPO	12.50 <sub>+9.11</sub>	10.10 <sub>+8.59</sub>	64.72 <sub>+40.82</sub>	57.15 <sub>+11.97</sub>	23.28 <sub>+6.30</sub>	26.90 <sub>+8.63</sub>	32.44
FlowRL	14.22 <sub>+10.83</sub>	9.58 <sub>+8.07</sub>	52.92 <sub>+29.02</sub>	66.20 <sub>+21.02</sub>	30.32 <sub>+13.34</sub>	34.47 <sub>+16.20</sub>	<b>34.62</b>

941  
942 Table 7: Ablation study on the effect of the  $\beta$  parameter in FlowRL. We report Avg@16 accuracy  
943 across six math reasoning benchmarks for different values of  $\beta$ .  
944

Models	AIME 2024	AIME 2025	AMC 2023	MATH-500	Minerva	Olympiad	Avg
$\beta = 5$	13.54	10.00	56.09	58.91	20.79	28.72	31.34
$\beta = 10$	14.79	10.20	59.53	64.30	25.27	32.39	34.41
$\beta = 15$	15.41	10.83	54.53	66.96	31.41	34.61	35.63
$\beta = 30$	15.00	10.83	50.62	69.02	30.03	35.03	35.09

## D HUMAN STUDY AND CROSS-DOMAIN EVALUATION

954  
955 **Human Evaluation.** We conduct a comprehensive human evaluation that demonstrates strong  
956 agreement with GPT-4o-mini assessments. We use the same rollouts from the GPT-4o-mini  
957 diversity experiment (Sec 6) to further validate diversity. As shown in Table 8, both evaluators  
958 independently identify FlowRL as the most diverse method and R++ as the least diverse, with GRPO  
959 and PPO showing intermediate diversity levels.

960 Human Instruction: As a human evaluator, assess the diversity of solutions for each problem by  
961 examining 16 solution attempts per method. Rate diversity on a 1-3 scale based on the following  
962 criteria:

- 963 • Score 1 (low diversity): 13+ responses use essentially identical approaches with only trivial dif-  
964 ferences in arithmetic, notation, or wording.
- 965 • Score 2 (moderate diversity): 7-12 responses use the most common approach, with 2-4 responses  
966 showing distinct alternative strategies.
- 967 • Score 3 (high diversity):  $\leq 6$  responses use the same method, with 4+ distinctly different solution  
968 strategies present.

969 **Other Domain Evaluation.** We conduct additional experiments on MMLU (Hendrycks et al.,  
970 2020) and GPQA (Rein et al., 2024) to demonstrate FlowRL’s effectiveness extends beyond math-  
971 ematical reasoning to other domains. We use Qwen-2.5-7B as the base model and follow the math  
972 training setup described in Sec 4. As shown in Table 9, FlowRL achieves the highest overall

972 scores on both benchmarks (72.13% on MMLU and 36.87% on GPQA). These results demonstrate  
 973 FlowRL’s strong generalization capability across different domains beyond the originally tested  
 974 mathematical reasoning tasks.

976  
 977 **Table 8: Human-evaluated diversity scores  
 978 (Mean  $\pm$  Std).**

Method	Score
R++	$1.10 \pm 0.20$
GRPO	$1.42 \pm 0.42$
PPO	$1.67 \pm 0.39$
FlowRL	$2.45 \pm 0.35$

976  
 977 **Table 9: MMLU and GPQA benchmark per-  
 978 formance.**

Method	MMLU	GPQA
R++	71.82	27.02
GRPO	71.87	33.08
PPO	72.10	33.84
FlowRL	72.13	36.87

## 985 986 E EXTENDED RELATED WORK AND COMPARISONS

989 Recent notable works have addressed similar challenges in large language model reinforcement  
 990 learning from different perspectives and across various domains. We provide a detailed comparison  
 991 below to highlight key distinctions and commonalities with existing methods.

993 **Length Normalization.** Dr. GRPO (Liu et al., 2025c) proposes an unbiased optimization method  
 994 that improves token efficiency by removing standard normalization terms from the advantage calcu-  
 995 lation and removing length terms from the loss objective, while focusing primarily on mathematical  
 996 reasoning improvements. SRPO (Zhang et al., 2025c) addresses length conflicts through a two-  
 997 stage training approach (math-first, then coding) and history resampling to filter zero-advantage  
 998 samples. GSPO (Zheng et al., 2025) conducts gradient analysis and applies length normalization  
 999 in the sequence-level importance ratio ( $s_i(\theta) = (\frac{\pi_\theta(y_i|x)}{\pi_{\theta_{\text{old}}}(y_i|x)})^{\frac{1}{|y_i|}}$ ) to avoid unstable training, partic-  
 1000 ularly crucial for MoE model training. FlowRL operates as a trajectory-level flow-balance objective  
 1001 that initially faced gradient explosion issues during long CoT reasoning. To overcome this chal-  
 1002 lenge, FlowRL integrates length normalization ( $\frac{1}{|y|} \log \pi_\theta(y|x)$ ) directly into the trajectory balance  
 1003 formulation, ensuring training stability and enabling effective scaling to extended CoT sequences.  
 1004 Unlike approaches requiring domain-specific training strategies, FlowRL’s unified formulation nat-  
 1005 urally handles variable sequence lengths through principled reward shaping within the flow-balance  
 1006 framework, achieving stable optimization across diverse reasoning tasks.

1008 **KL-Related Policy Optimization Methods.** Kimi-K1.5 (Team et al., 2025) employs on-  
 1009 policy sampling with KL regularization and uses empirical mean of sampled rewards ( $\bar{r} =$   
 1010  $\text{mean}(r(x, y_1, y^*), \dots, r(x, y_k, y^*))$ ) to approximate the normalizing constant  $Z$ . This objective has  
 1011 a closed form solution that introduces  $\log Z$ , where  $\gamma$  is a parameter controlling the degree of reg-  
 1012 ularization, maintaining the traditional reward maximization framework. IPO (Azar et al., 2024)  
 1013 addresses overfitting in preference-based learning by using identity mapping ( $\Psi = I$ ) to maintain  
 1014 effective KL regularization with deterministic preferences, targeting preference-based alignment  
 1015 problems. FlowRL differs by deriving its objective from reverse KL divergence minimization, shift-  
 1016 ing from reward maximization to reward distribution matching via flow balance. This approach  
 1017 employs a learnable partition function  $Z_\phi(x)$  parameterized by a 3-layer MLP and incorporates  
 1018 importance sampling for the entire trajectory balance objective. This approach provides both theo-  
 1019 retical rigor through generative flow networks and practical effectiveness across diverse reasoning  
 1020 tasks without requiring preference data or domain-specific training paradigms.

## 1021 1022 F IMPLEMENTATION OF PARTITION FUNCTION $Z_\phi$

1023 We detail the implementation of the partition function  $Z_\phi$ , covering theoretical foundations and  
 1024 practical aspects.

1026 From the flow perspective:  $Z_\phi$  measures the probability flow from the initial state  $S_0$ . Intuitively, it  
 1027 estimates the denominator—the sum of rewards across all possible paths—enabling conversion to a  
 1028 probability distribution via  $\frac{r(\mathbf{x}, \mathbf{y})}{Z_\phi(\mathbf{x})}$ .  
 1029

1030 From the implementation perspective: Since the input of  $Z_\phi$  corresponds to the initial state, we  
 1031 utilize the prompt representation from the language model. Specifically, we extract the hidden states  
 1032 from the final layer of the language model for all prompt tokens, and compute their mean to obtain  
 1033 a fixed-dimensional representation. This averaged hidden state vector serves as the input feature for  
 1034 computing the scalar partition function value  $\tilde{Z}_\phi(\mathbf{x})$ .

1035 We conduct comprehensive ablation studies examining: (1) MLP architecture depth (1/3/5 layers);  
 1036 (2) Removing  $Z_\phi$  entirely: to quantify how much  $Z_\phi$  contributes to the overall performance im-  
 1037 provement; (3) Replacing  $Z_\phi$  with a constant value: to assess whether adaptivity is necessary or a  
 1038 simple approximation suffices.

1039 The results demonstrate that the learnable partition function  $Z_\phi$  is essential for FlowRL’s per-  
 1040 formance. As shown in Table 10, varying MLP depth has minimal impact, with 3-layer MLP per-  
 1041 forming slightly better. Table 11 shows that removing  $Z_\phi$  causes significant drops (-5.62 on AIME  
 1042 2024, -6.25 on AIME 2025), while using a constant  $Z_\phi$  performs even worse (-7.91 and -8.75 re-  
 1043 spectively). These results confirm that  $Z_\phi$  is critical. Theoretically, it is essential for matching the  
 1044 reward distribution.

1045

1046 Table 10: MLP Architecture Depth.

$Z_\phi$ Arch.	AIME 2024	AIME 2025
1-layer MLP	12.79	8.12
3-layer MLP	15.41	10.83
5-layer MLP	10.49	6.77

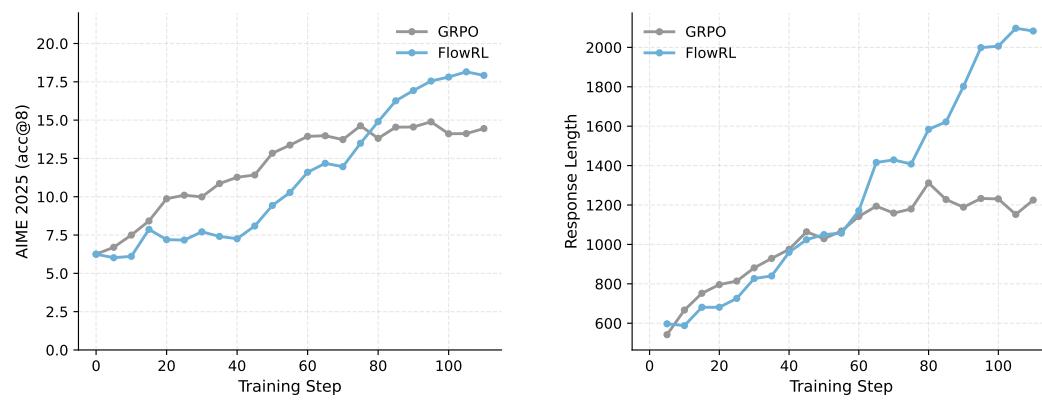
1047 Table 11: Partition Function  $Z_\phi$ .

Method	AIME 2024	AIME 2025
FlowRL	15.41	10.83
w/o $Z_\phi$	9.79	4.58
w/ constant $Z_\phi$	7.50	2.08

## G TRAINING ANALYSIS

1055 **Training Dynamics** We analyze model evolution during training by tracking AIME 2025 accuracy  
 1056 and response length. As shown in Figure 5, FlowRL gradually outperforms GRPO during training.

1057 FlowRL’s response length grows faster than GRPO, reaching approximately 2000 tokens by step  
 1058 100 compared to GRPO’s  $\sim$ 1200 tokens. Correspondingly, FlowRL achieves higher AIME 2025  
 1059 accuracy, with the performance gap widening as training progresses, particularly after step 75 where  
 1060 FlowRL begins to consistently outperform GRPO.



1075 Figure 5: Training dynamics on Qwen2.5-7B, including AIME 2025 Acc@8 (left) and response  
 1076 length (right).  
 1077

1078 **Reward Distribution Analysis.** We analyze reward distribution statistics during training on  
 1079 Qwen-2.5-32B. FlowRL maintains higher variance than GRPO, indicating exploration of diverse

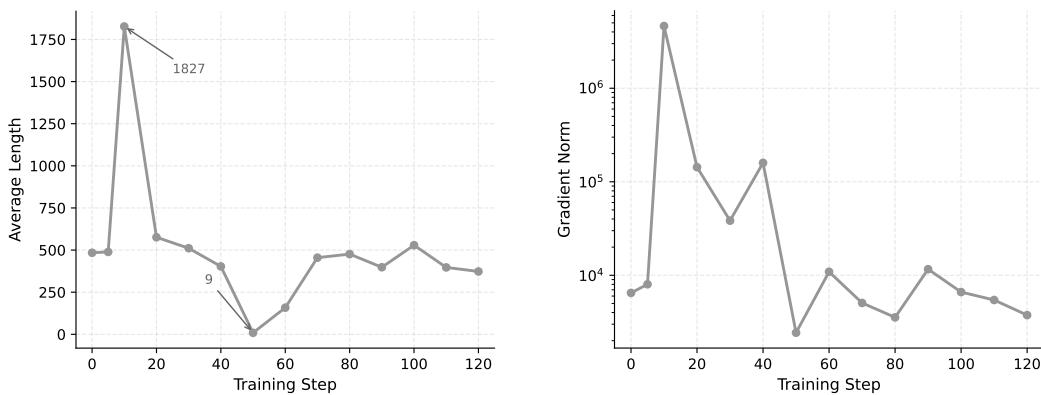
1080 solutions. Specifically, FlowRL achieves higher variance, aligning with flow matching theory that  
 1081 encourages exploration of multiple solution paths.  
 1082

1083 Table 12: Reward Distribution Statistics.  
 1084

Step	GRPO Std	FlowRL Std
0	0.1087	0.1087
50	0.1714	0.1341
100	0.0000	0.1165
150	0.0323	0.1664
200	0.1630	0.0730
245	0.0509	0.2341

1094 **Length Normalization Ablation.** We conduct an ablation study on the length normalization term  
 1095  $(1/|y|)$ . The results demonstrate that length normalization is essential for stable training.  
 1096

1097 Without it, training becomes highly unstable: at step 10, generation length explodes to 1827 tokens  
 1098 with gradient norm spiking to 4.6M; at step 50, length collapses to only 9 tokens, confirming that  
 1099 length normalization is critical for FlowRL’s stability.



1100  
 1101 Figure 6: **Ablation study on length normalization term  $(1/|y|)$ .** Left: average response length.  
 1102 Right: gradient norm (log scale). Without length normalization, training exhibits severe instability  
 1103 with length explosion/collapse and gradient spikes.  
 1104  
 1105  
 1106  
 1107  
 1108  
 1109  
 1110  
 1111  
 1112  
 1113  
 1114  
 1115  
 1116  
 1117  
 1118  
 1119

## H THE USE OF LARGE LANGUAGE MODELS

1120 LLMs (specifically GPT-4o-mini) are used as a judge to evaluate the diversity of solution approaches  
 1121 in our diversity analysis (Figure 4), following Yu et al. (2025a). All core research ideas, theoretical  
 1122 derivations, experimental design, and algorithmic innovations are developed by the authors without  
 1123 LLM assistance. The mathematical formulations and proofs are entirely the work of the human  
 1124 researchers. LLMs do not contribute to the fundamental conceptual development of FlowRL or the  
 1125 core insights about reward distribution matching via flow balance.  
 1126  
 1127  
 1128  
 1129  
 1130  
 1131  
 1132  
 1133

1134  
1135

## Diversity Evaluation Prompt

1136  
1137  
1138

**System:** You are evaluating the DIVERSITY of solution approaches for a mathematics competition problem. Focus on detecting even SUBTLE differences in methodology that indicate different problem-solving strategies.

1139

**PROBLEM:**

1140

{problem}

1141

**16 SOLUTION ATTEMPTS:**

1142

{formatted\_responses}

1143

**EVALUATION CRITERIA - Rate diversity from 1 to 5:**

1144

**Score 1 - Minimal Diversity:**

1145

- 14+ responses use essentially identical approaches
- Same mathematical setup, same variable choices, same solution path
- Only trivial differences (arithmetic, notation, wording)
- Indicates very low exploration/diversity in the generation process

1146

**Score 2 - Low Diversity:**

1147

- 11-13 responses use the same main approach
- 1-2 alternative approaches appear but are rare
- Minor variations within the dominant method (different substitutions, orderings)
- Some exploration but heavily biased toward one strategy

1148

**Score 3 - Moderate Diversity:**

1149

- 7-10 responses use the most common approach
- 2-3 distinct alternative approaches present
- Noticeable variation in problem setup or mathematical techniques
- Balanced mix showing reasonable exploration

1150

**Score 4 - High Diversity:**

1151

- 4-6 responses use the most common approach
- 3-4 distinct solution strategies well-represented
- Multiple mathematical techniques and problem framings
- Strong evidence of diverse exploration strategies

1152

**Score 5 - Maximum Diversity:**

1153

- No single approach dominates ( $\leq 3$  responses use same method)
- 4+ distinctly different solution strategies
- Wide variety of mathematical techniques and creative approaches
- Excellent exploration and generation diversity

1154

**IMPORTANT:** Focus on the DIVERSITY of the attempted approaches. Return ONLY a number from 1 to 5.

1155

1156

1157

1158

1159

1160

1161

1162

1163

1164

1165

1166

1167

1168

1169

1170

1171

1172

1173

1174

1175

1176

1177

1178

1179

1180

1181

1182

1183

1184

1185

1186

1187