

# MULTI-SCALE GROUP RELATIVE POLICY OPTIMIZATION FOR LARGE LANGUAGE MODELS

000  
001  
002  
003  
004  
005  
006  
007  
008  
009  
010  
011  
012  
013  
014  
015  
016  
017  
018  
019  
020  
021  
022  
023  
024  
025  
026  
027  
028  
029  
030  
031  
032  
033  
034  
035  
036  
037  
038  
039  
040  
041  
042  
043  
044  
045  
046  
047  
048  
049  
050  
051  
052  
053  
Anonymous authors  
Paper under double-blind review

## ABSTRACT

Reinforcement learning (RL) has become a cornerstone for improving the reasoning ability of large language models (LLMs). The current mainstream Group Relative Policy Optimization (GRPO) estimates advantage via relative comparisons within the full group of sampled responses. However, this single-scale, global comparison mechanism is inherently brittle, sensitive to the heterogeneity and stochasticity of reward distribution, leading to unstable training signals. Drawing inspiration from graph theory, where node importance is better captured through local substructures than global statistics, we propose *Multi-Scale Group Relative Policy Optimization* (MS-GRPO), a novel RL algorithm that generalizes GRPO by aggregating relative advantages computed across multiple response subgroups at varying scales (*e.g.*, pairwise, trios, etc.). Since the exhaustive enumeration of all meaningful subgroups grows combinatorially with group size, we further introduce a practical acceleration scheme that selects a small yet representative subset of subgroups via dilated scale sampling and diversity-aware subgroup selection. In addition, we provide a rigorous theoretical analysis, demonstrating that MS-GRPO can be interpreted as an adaptive correction of GRPO’s advantage controlled by the heterogeneity of reward distribution, and gracefully degenerates to GRPO when the reward distribution approaches homogeneity. Experiments demonstrate that MS-GRPO significantly outperforms GRPO on various tasks, for example, with improvements averaged over all evaluated models: +5.5 on AIME24 math reasoning, +4.6 on RiddleSense logical reasoning, +2.7 on LiveCodeBench programming challenges, +2.2 on MedQA medical reasoning, and +13.5 on HotpotQA with search engine.

## 1 INTRODUCTION

Large language models (LLMs) have demonstrated unprecedented capabilities in complex reasoning. A key driver behind this success is reinforcement learning (RL), which trains a policy to maximize a reward signal. As a cornerstone algorithm, Proximal Policy Optimization (PPO) (Schulman et al., 2017) suffers from training complexities and instabilities, largely due to its reliance on an online-trained value network for advantage estimation. Recently, Group Relative Policy Optimization (GRPO) (Shao et al., 2024b) emerged as an elegant alternative, which cleverly obviates the need for a learned value function. For each prompt, it samples a group of responses from the current policy model and uses their mean reward as an adaptive baseline. The advantage for each response is then computed by normalizing its reward relative to the group: subtracting the group’s mean reward and dividing by the group’s standard deviation.

Despite its conceptual simplicity and empirical success, GRPO’s advantage estimation mechanism suffers from a fundamental limitation: it performs a single-scale, global comparison across all responses in the group, thereby ignoring the rich, multi-granularity signals embedded in fine-grained, local comparisons. This global normalization is highly sensitive to reward distributional shifts and outlier responses, which are common in practice due to the stochasticity of LLM generation and the occasional instability of reward models. From a graph-theoretic perspective, treating each response as a node in a complete graph, GRPO’s approach is equivalent to characterizing each node solely by global graph statistics (mean and variance). However, decades of research in graph analysis have shown that such global characterizations are brittle and limited (Wu et al., 2020; Robinson et al., 2024; Segarra & Ribeiro, 2015; Valente et al., 2008): a node’s true importance or role is better captured by its participation in diverse local substructures (*e.g.*, motifs (Milo et al., 2002) and graphlets

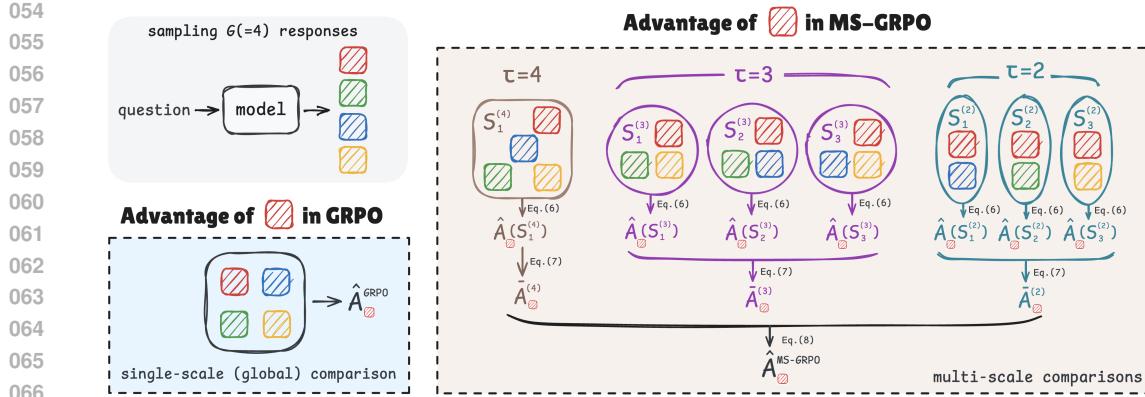


Figure 1: Comparison of the advantage estimation mechanism in MS-GRPO and GRPO. To make an intuitive comparison between them, we only illustrate the advantage of one response (denoted as a red square  $\square$ ). GRPO compares  $\square$  in the full response group, and normalizes its reward using the global mean and standard deviation of this full response group as its relative advantage. In contrast, our MS-GRPO compares  $\square$  in each valid subgroup containing it, and applies normalization within each comparison subgroup with varying scales (refer to Eq. (4)), and finally synthesizes these multi-view advantages to a holistic advantage via a hierarchical aggregation (refer to Eq. (5) and Eq. (6)).

(Pržulj, 2007)). Analogously, a response’s relative advantage should not be judged only against the entire group, but also against various local peer subsets (pairwise, trios, quartets, etc.), each offering a unique perspective on its advantage. Ignoring these multi-scale signals not only reduces the robustness of the advantage estimator but also forfeits valuable information related to the true and reliable value of a response.

In this paper, we propose Multi-Scale Group Relative Policy Optimization (MS-GRPO), a novel RL algorithm that generalizes GRPO by incorporating multi-scale relative comparisons into advantage estimation. Instead of GRPO’s single-scale advantage based on global comparison, MS-GRPO constructs a comprehensive advantage signal by aggregating relative advantages computed over all possible subgroups of responses, from pairwise comparisons up to the full group. To ensure statistical fairness, we introduce a hierarchical aggregation strategy that first averages advantages within each scale (*i.e.*, subgroup size) and then combines across scales with tunable weights, preventing larger or more numerous subgroup sizes from dominating the signal.

However, this multi-scale formulation, while conceptually powerful, entails a combinatorial explosion in the number of subgroups as the group size grows, resulting in more computational cost in practice. To address this challenge and enable scalable training, we design a practical **acceleration scheme** (refer to Sec. 3.2) that approximates the full subgroups via two complementary down-sampling strategies: (1) *Dilated Scale Sampling*, which sparsely selects a small set of representative scales across the granularity spectrum to reduce redundancy, and (2) *Diverse Group Sampling*, which, for each selected scale, chooses a compact yet maximally diverse subset of subgroups to preserve rich comparative information with minimal redundancy.

In the practical application of RL to LLMs, MS-GRPO demonstrates a notable superiority, stemming from both its solid theoretical foundation and its excellent experimental performance. From a **theoretical standpoint** as detailed in Appendix A, MS-GRPO modifies the advantage estimator of GRPO (Group Relative Policy Optimization) through an adaptive, heterogeneity-aware mechanism. When the reward distribution exhibits high heterogeneity<sup>1</sup>, MS-GRPO introduces additional corrections to effectively enhance the reliability of the advantage signal. At its core, MS-GRPO provides an extra advantage boost to samples with above-average rewards while penalizing those that fall below average, thereby improving the signal-to-noise ratio in unstable reward environments. It also places greater emphasis on the raw rewards of samples rather than their normalized relative rewards,

<sup>1</sup>The concept “**reward heterogeneity**” refers to the degree of non-uniformity in the reward distribution within the response group. High heterogeneity implies an uneven spread of rewards, often characterized by significant outliers or the emergence of distinct clusters (*e.g.*, high-reward *vs.* low-reward subgroups). This typically arises from randomness of the model’s performance or instability of the reward model.

108 which compensates for the diminished reliability of relative comparisons when the group’s reward  
 109 structure is unstable. When the reward distribution becomes homogeneous, MS-GRPO gracefully  
 110 degenerates to GRPO, ensuring the method provides robust, context-sensitive advantage estimation  
 111 when it matters most without introducing unnecessary corrections.

112 **Experimental results** further corroborate MS-GRPO’s superiority. Across multiple tasks aimed  
 113 at improving LLM reasoning, MS-GRPO consistently outperforms GRPO, delivering performance  
 114 boosts of +5.5 on AIME24 math reasoning, +4.6 on RiddleSense logical reasoning, +2.7 on Live-  
 115 CodeBench programming challenges, and +2.2 on MedQA medical reasoning. These improvements  
 116 are averaged over all evaluated model variants, including the Qwen2.5, LLaMA3.2 and DeepSeek-  
 117 R1-Distill-Qwen families with 1.5B, 3B, and 7B parameters. Even when LLMs are integrated with  
 118 search engines, MS-GRPO still significantly outperforms GRPO, achieving gains of +13.5 on Hot-  
 119 potQA benchmark. These experimental data powerfully prove that MS-GRPO can stably and sig-  
 120 nificantly improve LLM performance across various tasks and reward structures.

121

## 122 2 BACKGROUND

### 124 2.1 PRELIMINARY: GROUP RELATIVE POLICY OPTIMIZATION

126 Group Relative Policy Optimization (GRPO) (Shao et al., 2024b) circumvents the need for value  
 127 function approximation inherent in PPO (Schulman et al., 2017) by leveraging intra-group compar-  
 128 isons. Specifically, for each question  $q$ , GRPO samples a group of  $G$  responses,  $\mathcal{O} = \{o_1, \dots, o_G\}$   
 129 from the old policy model  $\pi_{\theta_{\text{old}}}$ . Each response  $o_i$  consists of a sequence of tokens  $(o_{i,1}, \dots, o_{i,|o_i|})$ ,  
 130 with a scalar reward  $r_{i,t} \in \mathbb{R}$  assigned to each token  $o_{i,t}$ . Let  $N = \sum_{i=1}^G |o_i|$  denote the total num-  
 131 ber of tokens across all responses in the group. The core of GRPO lies in its advantage estimation.  
 132 Taking response  $o_i$  as an example, it computes an advantage value  $\hat{A}_{i,t}^{\text{GRPO}}$  for each token in  $o_i$  by  
 133 normalizing its reward relative to the statistics of the entire group:

$$134 \hat{A}_{i,t}^{\text{GRPO}} = \frac{r_{i,t} - \mu_{\mathcal{O}}}{\sigma_{\mathcal{O}}}, \quad (1)$$

136 where  $\mu_{\mathcal{O}}$  and  $\sigma_{\mathcal{O}}$  are the mean and standard deviation of the rewards within the group  $\mathcal{O}$ , specifi-  
 137 cally,  $\mu_{\mathcal{O}} = \frac{1}{N} \sum_{i=1}^G \sum_{t=1}^{|o_i|} r_{i,t}$  and  $\sigma_{\mathcal{O}} = \sqrt{\frac{1}{N} \sum_{i=1}^G \sum_{t=1}^{|o_i|} (r_{i,t} - \mu_{\mathcal{O}})^2}$ . This advantage  $\hat{A}_{i,t}^{\text{GRPO}}$  is  
 138 subsequently used to optimize the policy model  $\pi_{\theta}$  by maximizing the following objective:  
 139

$$140 \mathcal{J}(\theta) = \mathbb{E}_{q \sim \mathcal{D}} \left\{ \frac{1}{G} \sum_{i=1}^G \frac{1}{|o_i|} \sum_{t=1}^{|o_i|} [\min(\lambda_{i,t} \hat{A}_{i,t}^{\text{GRPO}}, \text{clip}(\lambda_{i,t}, 1 - \epsilon, 1 + \epsilon) \hat{A}_{i,t}^{\text{GRPO}}) - \beta \mathbb{D}_{\text{KL}}(\pi_{\theta} || \pi_{\text{ref}})] \right\}, \quad (2)$$

143 where  $\mathcal{D}$  denotes the training dataset. The importance sampling ratio  $\lambda_{i,t} = \frac{\pi_{\theta}(o_{i,t}|q, o_{i,<t})}{\pi_{\theta_{\text{old}}}(o_{i,t}|q, o_{i,<t})}$  corrects  
 144 for the distributional shift between the behavior old policy model  $\pi_{\theta_{\text{old}}}$  that generated the responses  
 145 and the current policy model  $\pi_{\theta}$  being optimized. The  $\text{clip}(\cdot, 1 - \epsilon, 1 + \epsilon)$  operation serves to stabi-  
 146 lize training by constraining the policy update magnitude, where the clipping bounds are controlled  
 147 by the hyperparameter  $\epsilon$ .  
 148

### 149 2.2 MOTIVATION

151 Although GRPO provides an effective framework for policy alignment, its advantage estimation  
 152 mechanism is fundamentally limited by its reliance on a single, global comparison across the entire  
 153 group of  $G$  responses, thereby ignoring the rich signals embedded in more localized, fine-grained  
 154 comparisons. This limitation can be understood through the lens of graph theory. By conceptualizing  
 155 the responses as nodes in a complete graph ( $K_G$ ), GRPO is tantamount to characterizing each node  
 156 based solely on the graph’s global properties. However, a well-established principle in graph theory  
 157 is that characterizing a node solely by global graph properties can be misleading and brittle (Wu  
 158 et al., 2020; Robinson et al., 2024; Bringmann et al., 2019; Robins et al., 2007). Global properties  
 159 are highly sensitive to the presence of outlier nodes (Segarra & Ribeiro, 2015; Borgatti et al., 2006;  
 160 Albert et al., 2000; Cavallaro et al., 2024; Žnidaršič et al., 2018) (e.g., responses with exceptionally  
 161 high or low rewards) and the overall distribution of node attributes (Valente et al., 2008; Karimi  
 et al., 2018; Stoica et al., 2024; Salehzadeh-Yazdi & Hütt, 2025; Martin & Niemeyer, 2021) (e.g.,  
 reward distribution of a group of responses). Consequently, GRPO’s global comparison baseline is

162 easily skewed by the abnormal reward noise or specific reward distribution, obscuring a response’s  
 163 more nuanced role within its local neighborhood and yielding an advantage signal that lacks the  
 164 robustness and precision required for stable policy optimization.

165 To overcome this limitation of GRPO, we draw inspiration from a key concept in graph analysis:  
 166 characterizing a node by the ensemble of local subgraphs it participates in, which are often re-  
 167 ferred to as graphlets (Pržulj, 2007) or network motifs (Milo et al., 2002). In this field, a node’s  
 168 robust identity is defined not by its global position, but by this rich, multi-scale signature of its local  
 169 environment. Analogously, we propose that a more robust advantage signal can be derived by aggre-  
 170 gating a response’s relative performance across a diverse set of induced subgraphs of varying scales.  
 171 For instance, a global comparison might only reveal that a response is above average in the entire  
 172 group (analogous to a model’s average score on broad benchmarks), while our multi-scale approach  
 173 provides a richer comparison (analogous to a full performance breakdown, including its outstanding  
 174 performance on some key subfields). We name this paradigm *Multi-Scale Group Relative Policy*  
 175 *Optimization (MS-GRPO)* and detail its formulation in the following section.

### 176 3 METHOD

#### 177 3.1 MULTI-SCALE GROUP RELATIVE POLICY OPTIMIZATION

178 We propose Multi-Scale Group Relative Policy Optimization (MS-GRPO), whose core innovation  
 179 lies in a novel advantage estimation mechanism that produces a more robust and reliable advantage  
 180 by combining the advantages derived from response groups at varying scales. In contrast to GRPO,  
 181 which performs single-scale advantage estimation via global normalization over the full group of  
 182 sampled responses for a given question, our MS-GRPO first calculates a set of single-scale advan-  
 183 tages by independently normalizing rewards within every valid combination (subset) of responses,  
 184 ranging from the minimal pairwise combination to the full group. These single-scale advantages at  
 185 different scales are then fused via a hierarchical aggregation that explicitly balances the contribution  
 186 of each subset size, mitigating bias stemming from differing numbers of subsets at each size. The  
 187 resulting multi-scale advantage is used as the advantage signal in the policy optimization objective.  
 188

##### 189 3.1.1 MULTI-SCALE ADVANTAGE ESTIMATION

190 For a given question  $q$ , MS-GRPO samples a group of  $G$  responses  $\mathcal{O} = \{o_1, \dots, o_G\}$  from the old  
 191 policy model  $\pi_{\theta_{\text{old}}}$ . Each response  $o_i = (o_{i,1}, \dots, o_{i,|o_i|})$  is a token sequence, with corresponding  
 192 token-level rewards  $\{r_{i,t}\}_{t=1}^{|o_i|}$ , where  $r_{i,t} \in \mathbb{R}$  is the scalar reward assigned to token  $o_{i,t}$ . Our multi-  
 193 scale advantage estimation proceeds in the following three steps:

194 **(1) Group Construction.** Let’s define  $\mathbb{S}$  as the  $\tau_{\min}$ -power-set of the responses  $\mathcal{O}$ . Different from  
 195 the normal power set that contains all subsets of  $\mathcal{O}$ , the  $\tau_{\min}$ -power-set  $\mathbb{S}$  only considers all sub-  
 196 groups that contain at least  $\tau_{\min}$  responses. Formally, we have:

$$200 \quad \mathbb{S} = \{\mathbb{S}^{(\tau_{\min})}, \dots, \mathbb{S}^{(G)}\}, \quad \text{where } \mathbb{S}^{(\tau)} = \{\mathcal{S} \subseteq \mathcal{O} \mid |\mathcal{S}| = \tau\}, \quad (3)$$

201 where set  $\mathbb{S}^{(\tau)}$  is composed of all comparison subgroups with the same scale  $\tau$ . The hyperparam-  
 202 eter  $\tau_{\min}$  controlling the minimum comparison scale, with a default value of  $\tau_{\min} = 2$ , that is,  
 203 corresponding to pairwise comparison.

204 **(2) Advantage Estimation.** The core objective of this step is to capture each response’s relative  
 205 advantage within each subgroup. Specifically, for each response  $o_i$ , we compute its advantages  
 206 relative to peers in every comparison group containing it, *i.e.*,  $\forall \mathcal{S} \in \mathbb{S} \quad o_i \in \mathcal{S}$ . We define the advantage  
 207 assigned to token  $o_{i,t}$  within comparison group  $\mathcal{S}$  as

$$208 \quad \hat{A}_{i,t}(\mathcal{S}) = \frac{r_{i,t} - \mu_{\mathcal{S}}}{\sigma_{\mathcal{S}}}, \quad (4)$$

209 where  $\mu_{\mathcal{S}}$  and  $\sigma_{\mathcal{S}}$  are the mean and standard deviation of all tokens rewards with in the sub-group  
 210  $\mathcal{S}$ . After this step, each token  $o_{i,t}$  is associated with a collection of advantages  $\{\hat{A}_{i,t}(\mathcal{S})\}_{\mathcal{S} \in \mathbb{S}, o_i \in \mathcal{S}}$ ,  
 211 where each advantage quantifies the token’s relative performance within a specific comparison group  
 212  $\mathcal{S}$ , thereby establishing a multi-perspective basis for robust advantage estimation.

213 **(3) Hierarchical Aggregation.** This step synthesizes the multi-perspective advantages from Step 2  
 214 into a comprehensive advantage signal for each token that reflects its holistic relative merit against

216 the full response group  $\mathcal{O}$ . Specifically, for each token  $o_{i,t}$ , we aggregate the advantages  $\hat{A}_{i,t}(\mathcal{S})$   
 217 across all comparison groups  $\mathcal{S}$  containing  $o_i$ . However, simply averaging over all  $\hat{A}_{i,t}(\mathcal{S})$  would  
 218 introduce a statistical bias due to the unequal number of comparison groups per size. For response  
 219  $o_i$ , the number of comparison groups of size  $\tau$  containing  $o_i$  is  $\binom{G-1}{\tau-1}$ , since  $o_i$  is fixed to be included  
 220 and the remaining  $\tau-1$  responses are chosen from the other  $G-1$  responses in the full group  $\mathcal{O}$ .  
 221 This quantity peaks near  $\tau = G/2$ . As a result, medium-sized comparison groups would dominate  
 222 the aggregation, not because they provide higher-quality signals, but solely because they are more  
 223 numerous, leading to a biased advantage estimation. To avoid such bias, we perform a hierarchical  
 224 aggregation as follows: (1) For each token  $o_{i,t}$ , we compute the scale-specific averaged advantage  
 225 at each scale  $\tau \in \{\tau_{\min}, \dots, G\}$ :

$$\bar{A}_{i,t}^{(\tau)} = \frac{1}{\binom{G-1}{\tau-1}} \sum_{\substack{\mathcal{S} \in \mathcal{S} \\ |\mathcal{S}|=\tau, o_i \in \mathcal{S}}} \hat{A}_{i,t}(\mathcal{S}), \quad (5)$$

226 which averages all advantages from comparison groups of size  $\tau$  that contain  $o_i$ , yielding an aver-  
 227 aged advantage at a specific scale. (2) We combine these scale-specific advantages into a holistic  
 228 multi-scale advantage:

$$\hat{A}_{i,t}^{\text{MS-GRPO}} = \sum_{\tau=\tau_{\min}}^G w_\tau \cdot \bar{A}_{i,t}^{(\tau)}, \quad \text{where } w_\tau \geq 0 \text{ and } \sum_{\tau=\tau_{\min}}^G w_\tau = 1. \quad (6)$$

229 The weight coefficient  $w_\tau$  is a hyperparameter controlling the contribution of scale  $\tau$ . Uniform  
 230 weights ( $w_\tau = \frac{1}{G-\tau_{\min}+1}$ ) are used by default, treating all scales equally. Alternatively, manually-  
 231 designed weighting schemes can also be flexibly implemented, such as assigning higher weights to  
 232 larger scales (corresponding to larger comparison groups).

### 240 3.1.2 TRAINING OBJECTIVE

241 Building on our multi-scale advantage  $\hat{A}_{i,t}^{\text{MS-GRPO}}$  from Eq.(6), policy optimization proceeds by  
 242 adopting the GRPO-style objective defined in Eq.(2), with the sole modification that the original  
 243 advantage term  $\hat{A}_{i,t}^{\text{GRPO}}$  is replaced with our more reliable multi-scale advantage  $\hat{A}_{i,t}^{\text{MS-GRPO}}$ . This  
 244 substitution enables more reliable policy updates by leveraging multi-scale relative comparisons.

## 246 3.2 PRACTICAL ACCELERATION SCHEME

248 Although the multi-scale advantage estimation of MS-GRPO described in Sec. 3.1.1 is theoretically  
 249 sound, its computational complexity presents a practical scalability challenge. The total number of  
 250 valid comparison groups (*i.e.*, the size of  $\mathcal{S}$  defined in Eq.(3)) grows exponentially with the group  
 251 size  $G$ , rendering the computation of the multi-scale advantage computationally intractable for large  
 252 group size  $G$ .

253 To ensure the scalability of MS-GRPO, we introduce a practical acceleration scheme based on ap-  
 254 proximation via subsampling. The core idea is to compute an approximated multi-scale advantage  
 255 by operating not on the exhaustive set  $\mathcal{S}$ , but on a much smaller, representative subset  $\mathcal{C} \subset \mathcal{S}$ . The  
 256 construction of this subset  $\mathcal{C}$  is approached along two orthogonal dimensions to reduce redundancy.  
 257 First, we apply a **Dilated Scale Sampling** strategy to select a representative but sparse set of scales,  
 258 denoted as  $\mathbb{T}$ . Second, for each scale  $\tau \in \mathbb{T}$ , we employ a **Diverse Group Sampling** procedure to  
 259 select a concise yet informative subset of comparison groups, denoted as  $\mathcal{C}^{(\tau)}$ . The final, overall  
 260 representative subset  $\mathcal{C}$  is then formed by the union of these per-scale subsets  $\mathcal{C} = \cup_{\tau \in \mathbb{T}} \mathcal{C}^{(\tau)}$ . The  
 261 advantage estimation process (Steps (2) and (3) in Sec. 3.1.1) is then performed exclusively on this  
 262 condensed set  $\mathcal{C}$ . The specific mechanisms for each of these two sampling dimensions are detailed  
 263 in the following Sec. 3.2.1 and Sec. 3.2.2, respectively.

### 264 3.2.1 DILATED SCALE SAMPLING FOR SCALE SELECTION

266 In MS-GRPO, a *scale* refers to the size  $\tau = |\mathcal{C}|$  of a comparison group, which determines the gran-  
 267 ularity of the relative comparison. Small scales (*e.g.*,  $\tau = 2$ ) capture fine-grained, local comparisons,  
 268 while large scales (*e.g.*,  $\tau = G$ ) reflect coarse-grained, global comparisons. A full estimation that  
 269 considers all consecutive scales inherently introduces redundancy, as adjacent scales (*e.g.*,  $\tau$  and  
 $\tau+1$ ) yield highly correlated advantage signals.

Inspired by the classical dilated convolution (Yu & Koltun, 2015), we perform analogous dilated sampling over the scale dimension to ensure balanced coverage across the entire scale spectrum. Specifically, given a hyperparameter  $M$  representing the target number of scales, our *Dilated Scale Sampling* strategy partitions the full scale range  $[\tau_{\min}, G]$  into  $M$  non-overlapping intervals and selects one scale per interval. This process is formalized in two steps: **First**, we partition the integer range  $[\tau_{\min}, G]$  into  $M$  contiguous and non-overlapping intervals,  $\{I_j\}_{j=1}^M$ . The  $j$ -th interval,  $I_j = [s_j, e_j]$ , is formally defined by its start point  $s_j = \max(\tau_{\min}, e_{j-1} + 1)$  and end point  $e_j = G - (M - j) \cdot \lceil \frac{G - \tau_{\min} + 1}{M} \rceil$ . **Second**, we uniformly sample one scale  $\tau_j \sim \text{Uniform}(I_j)$  from each interval. This process yields the final set of selected scales  $\mathbb{T} = \{\tau_1, \dots, \tau_M\}$ , which forms the basis for the per-scale group sampling described next.

### 3.2.2 DIVERSE GROUP SAMPLING FOR PER-SCALE COMPARISON GROUP SELECTION

For each scale  $\tau \in \mathbb{T}$  selected above, the number of possible comparison groups is  $\binom{G}{\tau}$ , which can be prohibitively large and informationally redundant. For instance, two groups that differ by only one member provide highly correlated advantage signals, and averaging over them yields diminishing returns. Our goal, therefore, is to select a concise subset of comparison groups that maximizes internal diversity. To achieve this, we apply a conditional sampling strategy controlled by a budget hyperparameter  $K$ . Let  $\mathbb{S}^{(\tau)} = \{\mathcal{S} \subseteq \mathcal{O} \mid |\mathcal{S}| = \tau\}$  denote the set of all valid comparison groups at scale  $\tau$ , and  $N_\tau = |\mathbb{S}^{(\tau)}| = \binom{G}{\tau}$ . We define the representative subset  $\mathbb{C}^{(\tau)}$  as follows:

$$\mathbb{C}^{(\tau)} = \begin{cases} \mathbb{S}^{(\tau)} & \text{if } N_\tau \leq K, \\ K \text{ diverse comparison groups from } \mathbb{S}^{(\tau)} & \text{otherwise.} \end{cases} \quad (7)$$

In the latter case, we formalize the selection as a *diversity maximization optimization problem*. Specifically, we aim to find a subset  $\mathbb{C}^{(\tau)} \subset \mathbb{S}^{(\tau)}$  of size  $K$  that maximizes the sum of pairwise Jaccard distances  $d_J$  between its members:

$$\mathbb{C}^{(\tau)*} = \arg \max_{\mathbb{C}^{(\tau)} \subset \mathbb{S}^{(\tau)}, |\mathbb{C}^{(\tau)}| = K} \sum_{\mathcal{S}_a, \mathcal{S}_b \in \mathbb{C}^{(\tau)}, a \neq b} d_J(\mathcal{S}_a, \mathcal{S}_b), \quad (8)$$

where  $d_J(\mathcal{S}_a, \mathcal{S}_b) = 1 - \frac{|\mathcal{S}_a \cap \mathcal{S}_b|}{|\mathcal{S}_a \cup \mathcal{S}_b|}$  defines the divergence between any two groups  $\mathcal{S}_a, \mathcal{S}_b \in \mathbb{S}^{(\tau)}$ . As this problem is NP-hard, we design a fast, polynomial-time greedy algorithm that offers theoretical approximation guarantees (derived in Appendix E). Specifically, starting with an empty set  $\mathbb{C}^{(\tau)} = \emptyset$ , the algorithm iteratively constructs the set  $\mathbb{C}^{(\tau)}$  by repeatedly adding the candidate group  $\mathcal{S}^*$  from the remaining pool  $(\mathbb{S}^{(\tau)} \setminus \mathbb{C}^{(\tau)})$  that is most dissimilar to the groups already selected:

$$\mathcal{S}^* = \arg \max_{\mathcal{S} \in \mathbb{S}^{(\tau)} \setminus \mathbb{C}^{(\tau)}} \sum_{\mathcal{S}' \in \mathbb{C}^{(\tau)}} d_J(\mathcal{S}, \mathcal{S}'). \quad (9)$$

This greedy process terminates after  $K$  iterations, yielding the final representative set  $\mathbb{C}^{(\tau)}$  of size  $K$ . By construction, for each selected scale  $\tau \in \mathbb{T}$ , we have  $|\mathbb{C}^{(\tau)}| \leq K$ . Since  $|\mathbb{T}| = M$ , the total number of comparison groups used in the accelerated MS-GRPO is bounded by  $M \cdot K$ , ensuring computational tractability without sacrificing representativeness.

## 4 EXPERIMENTS

In this section, we first compare the performance of our MS-GRPO and GRPO on a wide range of tasks in Sec. 4.1, and then ablate each of our key designs in Sec. 4.2. The experimental settings are detailed in Appendix B.

### 4.1 MAIN RESULTS

**Math Reasoning.** As shown in Table 1, our MS-GRPO demonstrates comprehensive superiority on all math reasoning benchmarks, significantly and consistently outperforming GRPO across both the Qwen and LLaMA series models. For example, it achieves additional average accuracy gains of +4.1 for Qwen2.5-Math-7B and +5.0 for LLaMA3.2-3B-Instruct over GRPO across the five benchmarks. It is worth noting that as the model size increases (from 1.5B to 7B), the improvement brought by our

324 Table 1: Comparison with GRPO on five challenging math reasoning benchmarks.  
325

Model	AIME24	AMC23	MATH-500	MinervaMath	OlympiadBench	Avg.
Qwen2.5-Math-1.5B	13.3	30.0	36.6	19.1	22.6	24.3
+ GRPO	20.0	57.5	76.4	32.3	38.5	44.9
+ MS-GRPO (Ours)	26.7 (+6.7)	62.5 (+5.0)	80.0 (+3.6)	33.4 (+1.1)	40.5 (+2.0)	48.6 (+3.7)
Qwen2.5-Math-7B	13.3	40.0	53.6	17.2	17.4	28.3
+ GRPO	33.3	67.5	82.0	36.0	42.6	52.3
+ MS-GRPO (Ours)	40.0 (+6.7)	72.5 (+5.0)	82.8 (+0.8)	40.4 (+4.4)	46.4 (+3.8)	56.4 (+4.1)
Qwen2.5-Math-1.5B-Instruct	10.0	60.0	74.2	32.3	39.5	43.2
+ GRPO	16.7	62.5	76.4	30.8	40.1	45.3
+ MS-GRPO (Ours)	20.0 (+3.3)	62.5 (+0.0)	77.6 (+1.2)	33.0 (+2.2)	41.7 (+1.6)	47.0 (+1.7)
Qwen2.5-Math-7B-Instruct	13.3	70.0	81.2	36.0	45.6	49.2
+ GRPO	16.7	70.0	82.6	38.7	46.8	51.0
+ MS-GRPO (Ours)	23.3 (+6.6)	75.0 (+5.0)	83.5 (+0.9)	40.4 (+1.7)	47.8 (+1.0)	54.0 (+3.0)
LLaMA3.2-3B-Instruct	3.3	22.5	48.0	16.5	14.5	21.0
+ GRPO	13.3	27.5	57.2	20.9	21.6	28.1
+ MS-GRPO (Ours)	20.0 (+6.7)	40.0 (+12.5)	59.6 (+2.4)	22.7 (+1.8)	23.2 (+1.6)	33.1 (+5.0)
DeepSeek-R1-Distill-Qwen-1.5B	28.8	62.9	82.8	26.5	43.3	48.9
+ GRPO	30.0	67.5	83.8	29.7	47.0	51.6
+ MS-GRPO (Ours)	33.2 (+3.2)	75.0 (+7.5)	86.0 (+2.2)	31.1 (+1.4)	49.4 (+2.4)	54.9 (+3.3)

341 Table 2: Comparison with GRPO on five code generation benchmarks.  
342

Model	LiveCodeBench	HumanEval	HumanEval+	MBPP	MBPP+	Avg.
Qwen2.5-7B-Instruct	22.4	86.4	80.5	75.6	66.7	66.3
+ GRPO	28.6	87.8	84.0	80.4	68.9	70.0
+ MS-GRPO (Ours)	30.6 (+2.0)	88.5 (+0.7)	84.8 (+0.8)	82.5 (+2.1)	70.8 (+1.9)	71.4 (+1.4)
Qwen2.5-Coder-7B-Instruct	30.7	86.0	83.2	82.5	69.7	70.4
+ GRPO	32.6	87.1	83.3	83.7	70.1	71.4
+ MS-GRPO (Ours)	36.0 (+3.4)	88.0 (+0.9)	84.2 (+0.9)	85.9 (+2.2)	74.4 (+4.3)	73.7 (+2.3)

350 MS-GRPO also grows instead of narrows, which highlights the excellent scalability of our method.  
351 Furthermore, our method exhibits a more pronounced superiority on more difficult benchmarks,  
352 such as AIME24. These results provide strong evidence that, compared to GRPO, our proposed  
353 MS-GRPO can more effectively unlock the deep reasoning potential of LLMs for solving complex  
354 mathematical problems.

355 **Code Reasoning.** For code generation, MS-GRPO showcases similarly strong performance. As  
356 shown in Table 2, MS-GRPO delivers consistent performance gains over GRPO across all five code  
357 benchmarks. In particular, MS-GRPO achieves higher advantage of the code-specialized model,  
358 Qwen2.5-Coder-7B-Instruct, than that of the general model, Qwen2.5-7B-Instruct (+2.3 *vs.* +1.4).  
359 Similar with the findings in math reasoning, the gains of MS-GRPO relative to GPRO becomes  
360 more substantial on challenging benchmarks, for example, +3.4 on LiveCodeBench and +4.3 on  
361 MBPP+. This clearly indicates that MS-GRPO not only excels at general-purpose code tasks but  
362 also effectively enhances LLMs’ ability to tackle complex programming challenges.

363 **Logical Reasoning.** Logical reasoning, often assessed with puzzles, is a key indicator of the intel-  
364 ligence of LLMs. We evaluate on the well-known RiddleSense benchmark. The results in Table 3  
365 clearly demonstrate the significant superiority of MS-GRPO compared with GRPO. On the smaller  
366 1.5B level models, MS-GRPO surpasses GRPO by +5.3 and +5.9, respectively. On the larger 7B  
367 models, the advantage remains notable at +4.0 and +3.3. It shows that MS-GRPO can effectively  
368 unlock and enhance the logical reasoning capabilities of models, particularly for smaller models.

369 **Medical Reasoning.** In the highly specialized medical domain, MS-GRPO continues to demonstrate  
370 its superiority over GRPO. As shown in Table 4, on the authoritative MedQA benchmark (based on  
371 the US Medical Licensing Examination), our MS-GRPO consistently outperforms GRPO with gains  
372 of +2.0 to +2.5 across all baseline LLMs. This result validates the generalizability and effectiveness  
373 of MS-GRPO in specialized domains, showcasing its power to further enhance LLMs’ ability on  
374 tasks requiring deep, domain-specific knowledge.

375 **Question Answering with Search Engine.** We also compare MS-GRPO with GRPO for training  
376 LLMs integrated with a search engine. In this setting, the model can query a search engine to re-  
377 trieve relevant external knowledge to aid its reasoning process. Therefore, this training setting is  
378 particularly challenging since it requires enhancing not only the model’s core reasoning abilities but

378  
379  
380  
381 Table 3: Comparison with GRPO on logical rea-  
382 soning benchmark (RiddleSense).

Model	before training	+ GRPO	+ MS-GRPO (Ours)
Qwen2.5-1.5B	6.4	65.0	70.3 (+5.3)
Qwen2.5-7B	60.2	73.8	77.8 (+4.0)
Qwen2.5-1.5B-Instruct	36.5	64.3	70.2 (+5.9)
Qwen2.5-7B-Instruct	65.5	76.0	79.3 (+3.3)

386  
387  
388 Table 5: Comparison with GRPO on general and multi-hop QA tasks for LLMs with search engine.

Benchmarks	Qwen2.5-1.5B			Qwen2.5-3B		
	before training	+ GRPO	+ MS-GRPO (Ours)	before training	+ GRPO	+ MS-GRPO (Ours)
NQ	0.2	19.4	38.4 (+19.0)	2.3	40.6	44.1 (+3.5)
HotpotQA	0.4	18.3	30.8 (+12.5)	2.1	28.4	42.8 (+14.4)
Avg.	0.3	18.9	34.6 (+15.7)	2.2	34.5	43.5 (+9.0)

396 also its proficiency in interacting with the search engine. As shown in Table 5, MS-GRPO comprehensively and significantly outperforms GRPO on both general QA benchmark (NQ) and multi-hop 397 QA benchmark (HotpotQA), with the average bonus of +15.7 and +9.0 over GRPO for Qwen2.5- 398 1.5B and Qwen2.5-3B, respectively. Notably, MS-GRPO is exceptionally effective at unlocking the 399 potential of smaller models. Specifically, for Qwen2.5-1.5B, it delivers an impressive average 400 performance boost of +15.7, and achieves a massive gain of +19.0 on NQ benchmark, nearly doubling 401 the performance of GRPO. Furthermore, for Qwen2.5-3B, MS-GRPO shows larger improvements 402 on the more challenging HotpotQA than NQ, achieving remarkable gains of +12.5 and +14.4 over 403 GRPO. These results clearly indicate that MS-GRPO not only enhances a model’s intrinsic 404 reasoning ability but also significantly optimizes its performance when acting as an agent that interacts 405 with external tools like a search engine.

406  
407  
408 4.2 ABLATION STUDY

409 We ablate the efficacy of two key components in MS-GRPO: (1) the *Hierarchical Aggregation* 410 strategy for multi-scale advantages, and (2) the *Practical Acceleration Scheme* for scalable 411 computation. Ablation experiment on each term is conducted under two training settings: training 412 LLaMA3.2-3B-Instruct on math dataset, and training Qwen2.5-Coder-7B-Instruct on code dataset. We 413 report the average pass@1 across five benchmarks for math and code evaluation, respectively.

414 **Efficacy of Hierarchical Aggregation.** Hierarchical aggregation (detailed in Sec. 3.1) is a core 415 component of MS-GRPO that first averages advantages within each scale and then combines across 416 scales with weighted summation, making a fair contribution for all scales. We investigate two 417 aspects of this design:

418 • **Necessity:** To evaluate its necessity, we compare it against a naive aggregation that directly 419 averages all per-subgroup advantages without scale-wise awareness. As shown in Table 6a (line 1-2), 420 this naive manner underperforms hierarchical aggregation by 1.5 on math and 0.8 on code, 421 confirming that explicit scale-wise averaging is crucial to avoid bias from combinatorial imbalance.

422 • **Impact of weighting scheme  $\{w_\tau\}$ :** We compare three different weight coefficients  $\{w_\tau\}_\tau$  in 423 Eq. (6): (1) *uniform*: equal weight per scale, (2) *global-biased*: higher weight for larger scales, 424 emphasizing global comparisons. (3) *local-biased*: higher weight on smaller scales, emphasizing 425 local comparisons. In all cases, the weights are normalized to sum to one. As Table 6a (line 2- 426 4), uniform weighting achieves the best performance, while the local-biased weighting performs 427 worst. Notably, even the global-biased variant outperforms GRPO by +1.4 and +2.0, demon- 428 strating that incorporating any form of multi-scale signal can improve upon GRPO’s single-scale 429 advantage estimation.

430 **Efficacy of Acceleration Scheme.** Our acceleration scheme (detailed in Sec. 3.2) consists of two 431 orthogonal components: Dilated Scale Sampling and Diverse Group Sampling. We analyze their 432 individual and combined effects below:

432 Table 6: Ablation study on two key designs of MS-GRPO: (a) Hierarchical Aggregation (HA) and its  
 433 weighting scheme; (b) Practical Acceleration scheme with two components, Dilated Scale Sampling  
 434 (DSS) and Diverse Group Sampling (DGS). The impact of hyperparameters  $M$  and  $K$  is evaluated  
 435 in (c) and (d), respectively, where the reported scores follow the format *math / code* performance.  
 436 The light blue shallow highlighted our default configuration. For reference, GRPO achieves 28.1  
 437 (math) and 71.4 (code) in average pass@1 (%), with a training time of 184 seconds (s) per step.

(a)				(b)					
HA	weighting scheme	Math (Avg)	Code (Avg)	DSS	DGS	Math (Avg)	Code (Avg)	Time per step (Avg)	Comparison Group Number
✗	-	31.6	72.9	✗	✗	33.3	74.0	233s	247
✓	uniform	33.1	73.7	✓	✗	33.3	74.0	220s	93~155
✓	global-biased	32.7	73.4	✗	✓	33.1	73.8	203s	49
✓	local-biased	31.3	72.5	✓	✓	33.1	73.7	199s	25

(c)				(d)						
DSS sampling mode	2	4	6	no sampling	DGS sampling mode	4	8	16	32	no sampling
fixed	30.4 / 72.0	31.5 / 73.2	32.9 / 73.4	33.1 / 73.8	random	32.6 / 72.7	32.6 / 72.9	32.9 / 73.5	33.1 / 73.9	33.3 / 74.0
random	31.6 / 72.7	33.1 / 73.7	33.1 / 73.7	33.1 / 73.8	optimized	32.7 / 73.2	33.1 / 73.7	33.0 / 73.7	33.3 / 73.9	33.3 / 74.0

452 • **Impact of Dilated Scale Sampling (DSS):** DSS reduces redundancy across scales by selecting  
 453 only  $M$  representative scales from the full range  $[\tau_{\min}, G]$ . As shown in Table 6b, applying DSS  
 454 only reduces the number of comparison groups by nearly half while maintaining the performance  
 455 (line 1-2). A key design choice in DSS is the sampling strategy. Instead of deterministically  
 456 selecting scales at fixed intervals (e.g.,  $\tau = 2, 4, 6, 8$ ), we partition the scale range into  $M$  con-  
 457 tinuous intervals and randomly sample one scale per interval. As Table 6c shows, this random  
 458 sampling strategy consistently outperforms fixed-interval sampling across all  $M$  values. There-  
 459 fore, random sampling within each interval is the merit of maintaining the performance during  
 460 scale sparsification. We attribute this gain to the stochasticity introduced during training: ran-  
 461 dom scale selection ensures that all scales have a chance to be used over time, reducing human-  
 462 designed bias in fixed-interval sampling and improving generalization. Table 6c further studies  
 463 the sensitivity to hyperparameter  $M$ , the number of scales to be retained. Performance improves  
 464 as  $M$  increases, but with diminishing returns beyond  $M = 4$ . Given the trade-off between effi-  
 465 ciency and performance, we adopt  $M = 4$  as the default.

466 • **Impact of Diverse Group Sampling (DGS):** DGS controls the number of comparison groups  
 467 per scale via a budget  $K$ . As shown in Table 6b and Table 6d, disabling DGS (i.e., using all  
 468 subgroups) yields the best performance, while enabling DGS with  $K = 8$  slightly reduces scores  
 469 (from 33.3 to 33.1 on math, 74.0 to 73.7 on code), a minor drop given the drastic reduction in  
 470 the number of comparison groups (from 247 to 25). Crucially, the *diversity-aware selection* in  
 471 DGS is essential. Table 6d shows that the optimized strategy consistently outperforms random  
 472 sampling, especially at small  $K$ , indicating that maximizing subgroup diversity preserves more  
 473 informative signals under tight budgets  $K$ . Finally, we analyze the effect of  $K$ . Performance  
 474 increases with  $K$  and saturates around  $K = 16$  or 32. However, even with  $K = 8$ , our method  
 475 achieves 99.4% of the full performance (33.1 vs. 33.3 on math), thus, we choose  $K = 8$  as the  
 476 default, striking an optimal balance between efficiency and signal fidelity.

## 477 5 CONCLUSION

478 We propose Multi-Scale Group Relative Policy Optimization (MS-GRPO), which improves upon  
 479 GRPO by leveraging multi-scale comparisons across response subgroups to generate more robust  
 480 and reliable advantage signals. MS-GRPO mitigates the brittleness of global normalization under  
 481 reward heterogeneity and stochasticity. We further introduce a practical acceleration scheme to  
 482 ensure its efficiency and scalability.

483 **LLM-Usage Statement.** The authors used a large language model to assist with language polishing,  
 484 grammar correction, and typo identification in this paper. The ideas, methodology, experimental  
 485 design, and results presented are the sole work of the authors.

486 ETHICS STATEMENT  
487488 Based on the ICLR Code of Ethics, we confirm that our research adheres to its principles. The pri-  
489 mary contribution of our work is to improve the training effectiveness, efficiency, and generalization  
490 performance of large language models during the Reinforcement Learning from Human Feedback  
491 (RLHF) and fine-tuning stages. We believe these advancements will allow large models to better  
492 serve society. We acknowledge the importance of the responsible application of this technology.  
493 Our research does not involve the collection or use of any new personally identifiable information,  
494 and all experiments were conducted on publicly available datasets.495 REPRODUCIBILITY STATEMENT  
496497 We are fully committed to ensuring that our research is reproducible. To this end, we've included  
498 all the essential details required for independent verification, including comprehensive descriptions  
499 of our model architectures, training procedures, and all relevant hyperparameter settings in the main  
500 text and supplementary materials. For complete transparency and to facilitate future research, we  
501 plan to make the source code and all associated datasets publicly available as soon as this paper is  
502 officially accepted and published.503 REFERENCES  
504505

506 Meta AI. Llama 3.2 model card, 2024. URL [https://github.com/meta-llama/llama-models/blob/main/models/llama3\\_2/MODEL\\_CARD.md](https://github.com/meta-llama/llama-models/blob/main/models/llama3_2/MODEL_CARD.md).

507 AIME. American invitational mathematics examination, 2024. URL [https://artofproblemsolving.com/wiki/index.php/AIME\\_Problems\\_and\\_Solutions](https://artofproblemsolving.com/wiki/index.php/AIME_Problems_and_Solutions).

508 Réka Albert, Hawoong Jeong, and Albert-László Barabási. Error and attack tolerance of complex  
509 networks. *nature*, 406(6794):378–382, 2000.

510 AMC. American mathematics competitions, 2023. URL [https://artofproblemsolving.com/wiki/index.php/2023\\_AMC\\_12A](https://artofproblemsolving.com/wiki/index.php/2023_AMC_12A).

511 Jacob Austin, Augustus Odena, Maxwell Nye, Maarten Bosma, Henryk Michalewski, David Dohan,  
512 Ellen Jiang, Carrie Cai, Michael Terry, Quoc Le, et al. Program synthesis with large language  
513 models. *arXiv preprint arXiv:2108.07732*, 2021.

514 Maciej Besta, Nils Blach, Ales Kubicek, Robert Gerstenberger, Michal Podstawski, Lukas Gi-  
515 aninazzi, Joanna Gajda, Tomasz Lehmann, Hubert Niewiadomski, Piotr Nyczek, and Torsten  
516 Hoefer. Graph of thoughts: Solving elaborate problems with large language models. *Pro-  
517 ceedings of the AAAI Conference on Artificial Intelligence*, 38(16):17682–17690, March 2024.  
518 ISSN 2159-5399. doi: 10.1609/aaai.v38i16.29720. URL <http://dx.doi.org/10.1609/aaai.v38i16.29720>.

519 Stephen P Borgatti, Kathleen M Carley, and David Krackhardt. On the robustness of centrality  
520 measures under conditions of imperfect data. *Social networks*, 28(2):124–136, 2006.

521 Laura F Bringmann, Timon Elmer, Sacha Epskamp, Robert W Krause, David Schoch, Marieke  
522 Wicher, Johanna TW Wigman, and Evelien Snippe. What do centrality measures measure in  
523 psychological networks? *Journal of abnormal psychology*, 128(8):892, 2019.

524 Lucia Cavallaro, Pasquale De Meo, Giacomo Fiumara, and Antonio Liotta. On the sensitivity of  
525 centrality metrics. *Plos one*, 19(5):e0299255, 2024.

526 Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde De Oliveira Pinto, Jared  
527 Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, et al. Evaluating large  
528 language models trained on code. *arXiv preprint arXiv:2107.03374*, 2021.

529 Tianzhe Chu, Yuxiang Zhai, Jihan Yang, Shengbang Tong, Saining Xie, Dale Schuurmans, Quoc V  
530 Le, Sergey Levine, and Yi Ma. Sft memorizes, rl generalizes: A comparative study of foundation  
531 model post-training. *arXiv preprint arXiv:2501.17161*, 2025a.

540 Xiangxiang Chu, Hailang Huang, Xiao Zhang, Fei Wei, and Yong Wang. Gpg: A simple and strong  
 541 reinforcement learning baseline for model reasoning, 2025b. URL <https://arxiv.org/abs/2504.02546>.

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

546 Xidong Feng, Ziyu Wan, Muning Wen, Stephen Marcus McAleer, Ying Wen, Weinan Zhang, and  
 547 Jun Wang. Alphazero-like tree-search can guide large language model decoding and training,  
 548 2024. URL <https://arxiv.org/abs/2309.17179>.

549 Kanishk Gandhi, Ayush Chakravarthy, Anikait Singh, Nathan Lile, and Noah D Goodman. Cognitive  
 550 behaviors that enable self-improving reasoners, or, four habits of highly effective stars. *arXiv*  
 551 preprint *arXiv:2503.01307*, 2025.

552 Bofei Gao, Feifan Song, Zhe Yang, Zefan Cai, Yibo Miao, Qingxiu Dong, Lei Li, Chenghao Ma,  
 553 Liang Chen, Runxin Xu, et al. Omni-math: A universal olympiad level mathematic benchmark  
 554 for large language models. *arXiv preprint arXiv:2410.07985*, 2024.

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

558 Chaoqun He, Renjie Luo, Yuzhuo Bai, Shengding Hu, Zhen Leng Thai, Junhao Shen, Jinyi Hu,  
 559 Xu Han, Yujie Huang, Yuxiang Zhang, et al. Olympiadbench: A challenging benchmark for  
 560 promoting agi with olympiad-level bilingual multimodal scientific problems. *arXiv preprint*  
 561 *arXiv:2402.14008*, 2024.

562 Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song,  
 563 and Jacob Steinhardt. Measuring mathematical problem solving with the math dataset. *arXiv*  
 564 preprint *arXiv:2103.03874*, 2021.

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

567 Binyuan Hui, Jian Yang, Zeyu Cui, Jiaxi Yang, Dayiheng Liu, Lei Zhang, Tianyu Liu, Jiajun Zhang,  
 568 Bowen Yu, Keming Lu, et al. Qwen2. 5-coder technical report. *arXiv preprint arXiv:2409.12186*,  
 569 2024.

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

573 Di Jin, Eileen Pan, Nassim Oufattolle, Wei-Hung Weng, Hanyi Fang, and Peter Szolovits. What dis-  
 574 ease does this patient have? a large-scale open domain question answering dataset from medical  
 575 exams. *Applied Sciences*, 11(14):6421, 2021.

576 Fariba Karimi, Mathieu Génois, Claudia Wagner, Philipp Singer, and Markus Strohmaier. Ho-  
 577 mophily influences ranking of minorities in social networks. *Scientific reports*, 8(1):11077, 2018.

578 Amirhossein Kazemnejad, Milad Aghajohari, Eva Portelance, Alessandro Sordoni, Siva Reddy,  
 579 Aaron Courville, and Nicolas Le Roux. Vineppo: Unlocking rl potential for llm reasoning through  
 580 refined credit assignment. *arXiv preprint arXiv:2410.01679*, 2024.

581 Takeshi Kojima, Shixiang (Shane) Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwa-  
 582 sawa. Large language models are zero-shot reasoners. In S. Koyejo, S. Mohamed, A. Agarwal,  
 583 D. Belgrave, K. Cho, and A. Oh (eds.), *Advances in Neural Information Processing Systems*,  
 584 volume 35, pp. 22199–22213. Curran Associates, Inc., 2022. URL [https://proceedings.neurips.cc/paper\\_files/paper/2022/file/8bb0d291acd4acf06ef112099c16f326-Paper-Conference.pdf](https://proceedings.neurips.cc/paper_files/paper/2022/file/8bb0d291acd4acf06ef112099c16f326-Paper-Conference.pdf).

594 Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur Parikh, Chris  
 595 Alberti, Danielle Epstein, Illia Polosukhin, Jacob Devlin, Kenton Lee, et al. Natural questions: a  
 596 benchmark for question answering research. *Transactions of the Association for Computational  
 597 Linguistics*, 7:453–466, 2019.

598 Aitor Lewkowycz, Anders Andreassen, David Dohan, Ethan Dyer, Henryk Michalewski, Vinay Ra-  
 599 masesh, Ambrose Sloane, Cem Anil, Imanol Schlag, Theo Gutman-Solo, et al. Solving quantitative  
 600 reasoning problems with language models. *Advances in neural information processing systems*,  
 601 35:3843–3857, 2022.

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

605 Bill Yuchen Lin, Ziyi Wu, Yichi Yang, Dong-Ho Lee, and Xiang Ren. Riddlesense: Reasoning  
 606 about riddle questions featuring linguistic creativity and commonsense knowledge. *arXiv preprint  
 607 arXiv:2101.00376*, 2021.

608 Jiawei Liu, Chunqiu Steven Xia, Yuyao Wang, and Lingming Zhang. Is your code generated by chat-  
 609 gpt really correct? rigorous evaluation of large language models for code generation. *Advances  
 610 in Neural Information Processing Systems*, 36:21558–21572, 2023.

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

614 Michael Luo, Sijun Tan, Justin Wong, Xiaoxiang Shi, William Y. Tang, Manan Roongta, Colin Cai,  
 615 Jeffrey Luo, Li Erran Li, Raluca Ada Popa, and Ion Stoica. Deepscaler: Surpassing o1-preview  
 616 with a 1.5b model by scaling rl. <https://pretty-radio-b75.notion.site/DeepScaleR-Surpassing-O1-Preview-with-a-1-5B-Model-by-Scaling-RL-19681902c1468005bed8ca303013a4e2>, 2025. Not-  
 617 ion Blog.

618 Christoph Martin and Peter Niemeyer. On the impact of network size and average degree on the  
 619 robustness of centrality measures. *Network Science*, 9(S1):S61–S82, 2021.

620 Ron Milo, Shai Shen-Orr, Shalev Itzkovitz, Nadav Kashtan, Dmitri Chklovskii, and Uri Alon. Net-  
 621 work motifs: simple building blocks of complex networks. *Science*, 298(5594):824–827, 2002.

622 Yingqian Min, Zhipeng Chen, Jinhao Jiang, Jie Chen, Jia Deng, Yiwen Hu, Yiru Tang, Jiapeng  
 623 Wang, Xiaoxue Cheng, Huatong Song, et al. Imitate, explore, and self-improve: A reproduction  
 624 report on slow-thinking reasoning systems. *arXiv preprint arXiv:2412.09413*, 2024.

625 Niklas Muennighoff, Zitong Yang, Weijia Shi, Xiang Lisa Li, Li Fei-Fei, Hannaneh Hajishirzi, Luke  
 626 Zettlemoyer, Percy Liang, Emmanuel Candès, and Tatsunori Hashimoto. s1: Simple test-time  
 627 scaling. *arXiv preprint arXiv:2501.19393*, 2025.

628 George L Nemhauser, Laurence A Wolsey, and Marshall L Fisher. An analysis of approximations  
 629 for maximizing submodular set functions—i. *Mathematical programming*, 14(1):265–294, 1978.

630 OpenAI. Learning to reason with llms, 2024. URL <https://openai.com/index/learning-to-reason-with-llms/>.

631 OpenAI. Introducing openai o3 and o4-mini, 2025a. URL <https://openai.com/index/introducing-o3-and-o4-mini/>.

632 OpenAI. Openai o3-mini, 2025b. URL <https://openai.com/index/openai-o3-mini/>.

633 Nataša Pržulj. Biological network comparison using graphlet degree distribution. *Bioinformatics*,  
 634 23(2):e177–e183, 2007.

635 Laria Reynolds and Kyle McDonell. Prompt programming for large language models: Beyond  
 636 the few-shot paradigm. In *Extended Abstracts of the 2021 CHI Conference on Human Factors  
 637 in Computing Systems*, CHI EA ’21, New York, NY, USA, 2021. Association for Computing  
 638 Machinery. ISBN 9781450380959. doi: 10.1145/3411763.3451760. URL <https://doi.org/10.1145/3411763.3451760>.

648 Garry Robins, Pip Pattison, Yuval Kalish, and Dean Lusher. An introduction to exponential random  
 649 graph ( $p^*$ ) models for social networks. *Social networks*, 29(2):173–191, 2007.  
 650

651 Maria M Robinson, Isabella C DeStefano, Edward Vul, and Timothy F Brady. Local but not global  
 652 graph theoretic measures of semantic networks generalize across tasks. *Behavior Research Methods*,  
 653 56(6):5279–5308, 2024.

654 Ali Salehzadeh-Yazdi and Marc-Thorsten Hütt. Assessing the impact of sampling bias on node  
 655 centralities in synthetic and biological networks. *npj Systems Biology and Applications*, 11(1):47,  
 656 2025.

657

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

660 Santiago Segarra and Alejandro Ribeiro. Stability and continuity of centrality measures in weighted  
 661 graphs. *IEEE Transactions on Signal Processing*, 64(3):543–555, 2015.

662

663 Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang,  
 664 Mingchuan Zhang, YK Li, Y Wu, et al. Deepseekmath: Pushing the limits of mathematical  
 665 reasoning in open language models. *arXiv preprint arXiv:2402.03300*, 2024a.

666 Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang,  
 667 Mingchuan Zhang, YK Li, Yang Wu, et al. Deepseekmath: Pushing the limits of mathematical  
 668 reasoning in open language models. *arXiv preprint arXiv:2402.03300*, 2024b.

669

670 Guangming Sheng, Chi Zhang, Zilingfeng Ye, Xibin Wu, Wang Zhang, Ru Zhang, Yanghua Peng,  
 671 Haibin Lin, and Chuan Wu. Hybridflow: A flexible and efficient rlhf framework. In *Proceedings  
 672 of the Twentieth European Conference on Computer Systems*, pp. 1279–1297, 2025.

673 Ana-Andreea Stoica, Nelly Litvak, and Augustin Chaintreau. Fairness rising from the ranks: Hits  
 674 and pagerank on homophilic networks. In *Proceedings of the ACM Web Conference 2024*, pp.  
 675 2594–2602, 2024.

676

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

680 Trieu Trinh, Yuhuai Wu, Quoc Le, He He, and Thang Luong. Solving olympiad geometry without  
 681 human demonstrations. *Nature*, 2024. doi: 10.1038/s41586-023-06747-5.

682

683 Thomas W Valente, Kathryn Coronges, Cynthia Lakon, and Elizabeth Costenbader. How correlated  
 684 are network centrality measures? *Connections (Toronto, Ont.)*, 28(1):16, 2008.

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

688

689 Tailin Wu, Hongyu Ren, Pan Li, and Jure Leskovec. Graph information bottleneck. *Advances in  
 690 Neural Information Processing Systems*, 33:20437–20448, 2020.

691

692 Yunhui Xia, Wei Shen, Yan Wang, Jason Klein Liu, Huifeng Sun, Siyue Wu, Jian Hu, and Xiaolong  
 693 Xu. Leetcodedataset: A temporal dataset for robust evaluation and efficient training of code llms.  
 694 *arXiv preprint arXiv:2504.14655*, 2025.

695

696 Huajian Xin, Z. Z. Ren, Junxiao Song, Zhihong Shao, Wanjia Zhao, Haocheng Wang, Bo Liu,  
 697 Liyue Zhang, Xuan Lu, Qiushi Du, Wenjun Gao, Qihao Zhu, Dejian Yang, Zhibin Gou, Z. F.  
 698 Wu, Fuli Luo, and Chong Ruan. Deepseek-prover-v1.5: Harnessing proof assistant feedback for  
 699 reinforcement learning and monte-carlo tree search, 2024. URL <https://arxiv.org/abs/2408.08152>.

700

701 An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li,  
 702 Dayiheng Liu, Fei Huang, Haoran Wei, et al. Qwen2. 5 technical report. *arXiv preprint  
 703 arXiv:2412.15115*, 2024a.

702 An Yang, Beichen Zhang, Binyuan Hui, Bofei Gao, Bowen Yu, Chengpeng Li, Dayiheng Liu, Jian-  
 703 hong Tu, Jingren Zhou, Junyang Lin, et al. Qwen2. 5-math technical report: Toward mathematical  
 704 expert model via self-improvement. *arXiv preprint arXiv:2409.12122*, 2024b.

705 Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William W Cohen, Ruslan Salakhutdinov,  
 706 and Christopher D Manning. Hotpotqa: A dataset for diverse, explainable multi-hop question  
 707 answering. *arXiv preprint arXiv:1809.09600*, 2018.

708 Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Tom Griffiths, Yuan Cao, and Karthik  
 709 Narasimhan. Tree of thoughts: Deliberate problem solving with large language models. *Ad-*  
 710 *vances in neural information processing systems*, 36:11809–11822, 2023.

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

713 Fisher Yu and Vladlen Koltun. Multi-scale context aggregation by dilated convolutions. *arXiv*  
 714 *preprint arXiv:1511.07122*, 2015.

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

718 Eric Zelikman, Yuhuai Wu, Jesse Mu, and Noah Goodman. STar: Bootstrapping reasoning with  
 719 reasoning. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho (eds.), *Ad-*  
 720 *vances in Neural Information Processing Systems*, 2022. URL [https://openreview.net/](https://openreview.net/forum?id=_3ELRdg2sgI)  
 721 [forum?id=\\_3ELRdg2sgI](https://openreview.net/forum?id=_3ELRdg2sgI).

722 Anja Žnidaršič, Anuška Ferligoj, and Patrick Doreian. Stability of centrality measures in valued net-  
 723 works regarding different actor non-response treatments and macro-network structures. *Network*  
 724 *Science*, 6(1):1–33, 2018.

725

726

727

728

729

730

731

732

733

734

735

736

737

738

739

740

741

742

743

744

745

746

747

748

749

750

751

752

753

754

755

756 A ANALYSIS ON THE ADVANTAGE DIFFERENCE OF MS-GRPO AND GRPO  
757

758 To further understand the distinction between the advantage estimators of MS-GRPO and GRPO,  
759 we derive an analytical expression for their difference, denoted as  $\Delta\hat{A}_{i,t}$ . The complete derivation  
760 is provided in Appendix D.1, yielding the following formula (where  $\text{Var}(\cdot)$  denotes the variance):  
761

$$762 \Delta\hat{A}_{i,t} = \hat{A}_{i,t}^{\text{MS-GRPO}} - \hat{A}_{i,t}^{\text{GRPO}} \\ 763 \\ 764 \approx \underbrace{(r_{i,t} - \mu_{\mathcal{O}})}_{\text{global-relative reward}} \underbrace{\left( \frac{1}{\mathbb{E}_{\mathcal{S}}[\sigma_{\mathcal{S}}]} - \frac{1}{\sigma_{\mathcal{O}}} \right)}_{\substack{\text{Correction Term 1} \\ \text{scaling factor 1}}} + \underbrace{r_{i,t}}_{\text{reward}} \underbrace{\frac{\text{Var}_{\mathcal{S}}(\sigma_{\mathcal{S}})}{(\mathbb{E}_{\mathcal{S}}[\sigma_{\mathcal{S}}])^3}}_{\substack{\text{Correction Term 2} \\ \text{scaling factor 2}}} \quad (10) \\ 765 \\ 766 \\ 767$$

768 This difference can be interpreted as an additive correction bias that MS-GRPO applies to the  
769 GRPO’s advantage. It consists of two distinct correction terms: the first scales the *global-relative*  
770 *reward* ( $r_{i,t} - \mu_{\mathcal{O}}$ ), while the second scales the *raw reward* ( $r_{i,t}$ ).  
771

772 **Rationale of Correction Term 1.** The first correction term can be understood as using the difference  
773 between each token’s reward  $r_{i,t}$  and the global average reward  $\mu_{\mathcal{O}}$  to correct its relative advantage.  
774 Since “scaling factor 1 ( $1/\mathbb{E}_{\mathcal{S}}[\sigma_{\mathcal{S}}] - 1/\sigma_{\mathcal{O}}$ )” is non-negative (proven in Appendix Proof 1), the sign  
775 of “correction term 1” is determined solely by whether the token’s reward exceeds or falls below the  
776 global average reward. Among all tokens in the G responses, it grants an additional advantage bonus  
777 to tokens above the global reward level, while imposing an additional advantage penalty to tokens  
778 below the global reward level. In addition to being affected by the absolute deviation  $|r_{i,t} - \mu_{\mathcal{O}}|$ , the  
779 magnitude of these additional advantage bonus or penalties also increases with the heterogeneity of  
780 the reward distribution, since the magnitude of the “scaling factor 1” increases with the heterogeneity  
781 of the reward distribution (proven in Appendix Proof 2). This is reasonable because if the model’s  
782 answer quality to the current question is unstable (high heterogeneity), a response that significantly  
783 outperforms the average is more notable and thus deserves a larger advantage boost.  
784

785 **Rationale of Correction Term 2.** The second correction term can be understood as using the scaled  
786 reward value of each token to correct its relative advantage. Since “scaling factor 2” is non-negative  
787 (proven in Appendix Proof 3), the sign of the second correction term is determined by the positive  
788 or negative value of the raw reward  $r_{i,t}$ . Among all tokens in the G responses, this term assigns  
789 a higher additional advantage boost for higher-quality tokens (corresponding to higher rewards).  
790 In addition to being affected by the reward value of the token, the magnitude of this correction  
791 term is also affected by the reward heterogeneity. The greater the heterogeneity, the larger the  
792 magnitude of the second correction term, because the magnitude of “scaling factor 2” increases  
793 with the heterogeneity of the reward distribution (proven in Appendix Proof 4). This is reasonable  
794 because if the reward distribution is highly heterogeneous (e.g., caused by the instability of model  
795 performance or reward model), the reward groups for advantage estimation will also be unstable,  
796 and the reliability of relative comparisons will be weakened. Therefore, in this highly uncertain  
797 comparison environment, it is reasonable to place slightly more trust in a token’s absolute superiority  
798 (reflected by its raw reward  $r_{i,t}$ ), as the second correction term does.  
799

800 **Overall Interpretation: MS-GRPO as an Adaptive Advantage Correction.** In summary, our  
801 analysis reveals that the advantage difference  $\Delta\hat{A}_{i,t}$  acts as an adaptive, heterogeneity-aware cor-  
802 rection that MS-GRPO applies to the GRPO’s advantage. The behavior of this correction is dictated  
803 by the structure of the reward distribution:  
804

- 805 • When the reward distribution is nearly *homogeneous*, the two variance related to the reward  
806 heterogeneity ( $\text{Var}_{\mathcal{S}}(\mu_{\mathcal{S}})$  and  $\text{Var}_{\mathcal{S}}(\sigma_{\mathcal{S}})$ ) approach zero. Consequently, both “scaling factor 1”  
807 and “scaling factor 2” also approach zero, causing the entire advantage difference  $\Delta\hat{A}_{i,t}$  to  
808 vanish. In this simple case, MS-GRPO gracefully degenerates to GRPO, demonstrating that  
809 the correction is adaptively applied only when necessary.
- 810 • Conversely, when the reward distribution is highly *heterogeneous*, the correction term becomes  
811 more significant. The correction operates through two synergistic mechanisms. “Correction  
812 Term 1” grants an additional advantage boost to tokens with above-average rewards (since out-  
813 performing the average is more significant and notable in an unstable sampling), and imposes  
814 an additional advantage penalty on those below average. Simultaneously, “Correction Term  
815 2” places greater trust in a token’s raw reward when the high heterogeneity makes the context  
816 more uncertain.

810 for relative comparisons unstable. Both terms, therefore, leverage the rich, fine-grained information  
 811 embedded in the group’s reward heterogeneity to produce a more nuanced and reliable  
 812 advantage signal. This stands in sharp contrast to GRPO’s single-scale, global estimation,  
 813 which completely overlooks the crucial structural information.

814 The combined effect results in a sophisticated correction that not only accounts for a token’s relative  
 815 standing but also the reliability of the context in which that standing is measured, leading to a more  
 816 robust and nuanced advantage signal.

## 818 B EXPERIMENTAL SETTINGS

821 **Implementation Details.** Our implementation is built upon the `VeRL` (Sheng et al., 2025) library.  
 822 For the hierarchical aggregation in Eq.(6), we adopt uniform weights across all scales (*i.e.*,  $w_\tau =$   
 823  $1/(G - \tau_{\min} + 1)$ ) by default, because we empirically find that this simple setting consistently  
 824 yields strong performance. In fact, linearly assigning higher weights to larger scales underperforms  
 825 uniform weighting, as verified in our ablation study (see Table 4.2 (a)). For the two hyperparameters  
 826 introduced in our acceleration scheme, we set  $M = 4$  and  $K = 8$  by default in our experiments.  
 827 These values strike a favorable trade-off between computational efficiency and advantage estimation  
 828 fidelity across all tasks, and we ablate the effect of them in Table 6c and Table 6d.

829 **Training Settings.** For each training data, we sample  $G = 8$  responses from the current policy with  
 830 temperature 1.0. The maximum output sequence length is set as 2048 tokens. We use a constant  
 831 learning rate of  $1 \times 10^{-6}$ , weight decay of 0.01, and the AdamW optimizer. The prompt batch  
 832 size is set to 256. Training proceeds for 500 steps for math reasoning, code reasoning, and LLM  
 833 with search engine tasks. While for logical and medical reasoning, we train for 300 steps due to  
 834 their faster convergence. The KL penalty coefficient in the objective function of GRPO is set to  
 835 0.0001, and we set the entropy coefficient as 0.001 to encourage exploration. All the rewards are  
 836 derived from rule-based outcome verifiers. For code reasoning, we set the pass rate over all test cases  
 837 ( $\frac{\#\text{passed test cases}}{\#\text{total test cases}}$ ) as the reward. For other tasks, a binary reward is assigned based on the correctness  
 838 of the final answer: the reward is set to 1 if the final answer is correct, otherwise, the reward is set  
 839 to 0.

840 **Training Dataset.** The training datasets we used are listed as follows: (1) *Math Reasoning*: We  
 841 train the DeepSeek-R1-Distill-Qwen-1.5B model on a combination of AIME (1984-2023), AMC  
 842 problems (prior to 2023), and Omni-MATH (Gao et al., 2024) datasets. Other models are trained  
 843 on 8K problems from the MATH (Hendrycks et al., 2021) dataset with difficulty levels between 3  
 844 and 5. (2) *Code Reasoning*: All models are trained on 2K programming problems from the Leet-  
 845 CodeDataset (Xia et al., 2025). (3) *Medical Reasoning*: All models are trained on 10K English  
 846 questions from the training split of MedQA (USMLE) (Jin et al., 2021). (4) *Logical Reasoning*: All  
 847 models are trained on 3K examples from the training split of RiddleSense (Lin et al., 2021). (5)  
 848 *QA with Search Engine*: All models are trained on 90K questions from the training set of HotPotQA  
 (Yang et al., 2018).

849 **Evaluation Protocol.** We employ greedy decoding for all test data and report the `pass@1` metric  
 850 across the following benchmarks: (1) For *Math Reasoning*: we evaluate the mathematics problem-  
 851 solving ability on five widely-used challenging benchmarks, including AIME24 (AIME, 2024),  
 852 AMC23 (AMC, 2023), MATH-500 (Hendrycks et al., 2021), MinervaMath (Lewkowycz et al.,  
 853 2022), and OlympiadBench (He et al., 2024). (2) For *Code Reasoning*: We evaluate the complex  
 854 programming ability on five code generation benchmarks, including MBPP (Austin et al., 2021),  
 855 MBPP+ (Liu et al., 2023), HumanEval (Chen et al., 2021), HumanEval+ (Liu et al., 2023), and  
 856 LiveCodeBench (Jain et al., 2024). (3) For *Logical Reasoning*: We evaluate the LLMs’ ability to  
 857 solve puzzles on the well-known RiddleSense benchmark. (4) For *Medical Reasoning*: The medical  
 858 reasoning evaluation is conducted using the test split of MedQA (USMLE) (Jin et al., 2021). (5)  
 859 *QA with Search Engine*: We evaluate on two popular QA benchmarks, one Natural Questions (NQ)  
 860 (Kwiatkowski et al., 2019) designed for general question answering, and another is HotPotQA (Yang  
 861 et al., 2018), which qualifies the multi-hop reasoning ability.

862 **Baseline Models.** We conduct experiments across a diverse set of popular LLMs, including both the  
 863 *Qwen2.5* (Yang et al., 2024a) and *LLaMA3.2* (AI, 2024) families with varying sizes (1.5B, 3B, and  
 7B parameters). For math and code reasoning, we additionally experiment with specialized models,  
 864 *Qwen2.5-Math* (Yang et al., 2024b) and *Qwen2.5-Coder* (Hui et al., 2024), respectively. Beyond

them, we also consider a LongCoT instruction-tuned model, *DeepSeek-R1-Distill-Qwen-1.5B* (Guo et al., 2025), to examine the effectiveness of our method on the model with the recent advanced long chain-of-thought reasoning paradigm. Both base LLMs and instruction-tuned LLMs are considered to provide comprehensive comparisons on two common RL settings: (1) Directly RL starting from a base (pretrained-only) model, *i.e.*, Zero-RL, and (2) Applying RL for an instruction-tuned (SFT) model.

## C RELATED WORK

**LLM Reasoning.** Recent research (OpenAI, 2024; 2025a;b; Guo et al., 2025) has demonstrated that large language models can achieve significant performance gains by incorporating step-by-step reasoning. Most of these approaches rely on prompting to guide the model into generating explicit and sequential reasoning paths (Wei et al., 2022; Yao et al., 2023; Besta et al., 2024). This includes well-known methods such as Chain-of-Thought (CoT) prompting (Wei et al., 2022; Kojima et al., 2022; Reynolds & McDonell, 2021; Zelikman et al., 2022; Ye et al., 2025), Tree-of-Thought (Yao et al., 2023), and Graph of Thoughts (Besta et al., 2024). Additionally, some works integrate more sophisticated search algorithms (Feng et al., 2024; Xin et al., 2024; Trinh et al., 2024) with the reasoning process, for instance, by using Monte Carlo Tree Search (Trinh et al., 2024). However, reasoning capabilities triggered solely by prompting can be unstable and may not fully unlock a model’s true potential. To address this limitation, other research has proposed training-based methods to further strengthen a model’s step-by-step reasoning ability. These efforts often involve creating datasets (Muennighoff et al., 2025; Min et al., 2024; Luo et al., 2025) with reasoning annotations or utilizing bootstrapping self-training techniques (Zelikman et al., 2022). To improve reasoning performance, some studies have trained verifiers to check the validity of intermediate steps (Lightman et al., 2023; Cobbe et al., 2021). The use of these verifiers has been shown to significantly boost a model’s overall reasoning ability.

**Reinforcement Learning for LLMs.** A key driver of recent advancements in large language models (LLMs) (OpenAI, 2024; 2025a;b; Guo et al., 2025; Team et al., 2025) has been Reinforcement Fine-Tuning (RFT) (Schulman et al., 2017; Shao et al., 2024a; Guo et al., 2025; Yu et al., 2025; Hu, 2025), a technique that refines model behavior through reward-guided optimization. This approach fundamentally differs from Supervised Fine-Tuning (SFT) (Muennighoff et al., 2025; Min et al., 2024; Luo et al., 2025), which aligns model outputs with labeled responses, by using reinforcement learning to adapt models based on feedback signals. Typically, RFT relies on reinforcement learning algorithms like Proximal Policy Optimization (PPO) (Schulman et al., 2017) paired with rule-based reward functions. For example, DeepSeek-R1 (Guo et al., 2025) utilized Group Relative Policy Optimization (GRPO) (Shao et al., 2024a) with binary rewards to indicate the correctness of answers in tasks such as mathematics (AIME, 2024) and coding (Jain et al., 2024), achieving impressive results. A number of studies suggest that RFT not only enhances cognitive abilities like reflection and self-correction (Gandhi et al., 2025; Guo et al., 2025) but also improves generalization across various tasks (Chu et al., 2025a). Current RFT research is largely centered on algorithmic improvements. For instance, VinePPO (Kazemnejad et al., 2024) was designed to address the limitations of PPO’s value networks in complex reasoning tasks by introducing unbiased Monte Carlo estimates for better credit assignment, leading to gains in both efficiency and performance over PPO baselines. Similarly, Liu et al. (2025) analyzed the training pipeline of DeepSeek-R1-Zero (Guo et al., 2025), identifying biases in GRPO and proposing Dr.GRPO to enhance both token efficiency and final performance. Other efforts have aimed to simplify the GRPO algorithm, for example, by removing the KL-divergence term to produce more robust empirical results (Yu et al., 2025; Chu et al., 2025b).

## D ANALYSIS ON THE ADVANTAGE ESTIMATION DIFFERENCE BETWEEN MS-GRPO AND GRPO

In this section, we first provide a detailed derivation for the advantage difference  $\Delta \hat{A}_{i,t}$  between Multi-Scale Group Relative Policy Optimization (MS-GRPO) and Group Relative Policy Optimization (GRPO) in Sec. D.1. Then, we prove the properties of two critical terms of  $\Delta \hat{A}_{i,t}$  in Sec. D.2 and Sec. D.3.

918 D.1 DERIVATION OF THE APPROXIMATE ADVANTAGE DIFFERENCE  
919

920 Let  $\hat{A}_{i,t}^{\text{MS-GRPO}}$  and  $\hat{A}_{i,t}^{\text{GRPO}}$  denote the advantage estimates from MS-GRPO and GRPO for  $o_{i,t}$  (the  
921  $t$ -th token of  $i$ -th response) with reward  $r_{i,t}$ , respectively. The difference is defined as:

$$922 \quad \Delta \hat{A}_{i,t} = \hat{A}_{i,t}^{\text{MS-GRPO}} - \hat{A}_{i,t}^{\text{GRPO}} \quad (11)$$

923 For clarity in the following formulas, we use the simplified notation  $\mathbb{E}_{\mathcal{S}}[\cdot]$  to represent the hierar-  
924 chical aggregation of MS-GRPO. Substituting the definitions of advantage in Eq. (6) and Eq. (1), we  
925 can expand the expression as follows:

$$\begin{aligned} 927 \quad \Delta \hat{A}_{i,t} &= \mathbb{E}_{\mathcal{S}} \left[ \frac{r_{i,t} - \mu_{\mathcal{S}}}{\sigma_{\mathcal{S}}} \right] - \frac{r_{i,t} - \mu_{\mathcal{O}}}{\sigma_{\mathcal{O}}} \\ 928 \\ 929 &= \mathbb{E}_{\mathcal{S}} \left[ \frac{r_{i,t}}{\sigma_{\mathcal{S}}} \right] - \mathbb{E}_{\mathcal{S}} \left[ \frac{\mu_{\mathcal{S}}}{\sigma_{\mathcal{S}}} \right] - \frac{r_{i,t}}{\sigma_{\mathcal{O}}} + \frac{\mu_{\mathcal{O}}}{\sigma_{\mathcal{O}}} \\ 930 \\ 931 &= r_{i,t} \mathbb{E}_{\mathcal{S}} \left[ \frac{1}{\sigma_{\mathcal{S}}} \right] - \mathbb{E}_{\mathcal{S}} \left[ \frac{\mu_{\mathcal{S}}}{\sigma_{\mathcal{S}}} \right] - \frac{r_{i,t}}{\sigma_{\mathcal{O}}} + \frac{\mu_{\mathcal{O}}}{\sigma_{\mathcal{O}}} \end{aligned} \quad (12)$$

934 To make this expression tractable, we further derive the expressions for the two expectation terms  
935  $\mathbb{E}_{\mathcal{S}} \left[ \frac{1}{\sigma_{\mathcal{S}}} \right]$  and  $\mathbb{E}_{\mathcal{S}} \left[ \frac{\mu_{\mathcal{S}}}{\sigma_{\mathcal{S}}} \right]$ :

- 937 • **Derivation for  $\mathbb{E}_{\mathcal{S}} \left[ \frac{1}{\sigma_{\mathcal{S}}} \right]$ :** We apply the second-order Taylor expansion on function  $f(\sigma_{\mathcal{S}}) =$   
938  $1/\sigma_{\mathcal{S}}$  around the mean  $\mathbb{E}_{\mathcal{S}}[\sigma_{\mathcal{S}}]$ :

$$940 \quad f(\sigma_{\mathcal{S}}) \approx f(\mathbb{E}_{\mathcal{S}}[\sigma_{\mathcal{S}}]) + (\sigma_{\mathcal{S}} - \mathbb{E}_{\mathcal{S}}[\sigma_{\mathcal{S}}]) f'(\mathbb{E}_{\mathcal{S}}[\sigma_{\mathcal{S}}]) + \frac{1}{2} (\sigma_{\mathcal{S}} - \mathbb{E}_{\mathcal{S}}[\sigma_{\mathcal{S}}])^2 f''(\mathbb{E}_{\mathcal{S}}[\sigma_{\mathcal{S}}]) \quad (13)$$

941 Taking the expectation of both sides, the linear term vanishes since  $\mathbb{E}_{\mathcal{S}}[\sigma_{\mathcal{S}} - \mathbb{E}_{\mathcal{S}}[\sigma_{\mathcal{S}}]] = 0$ , and  
942 we get the approximation for the expectation:

$$\begin{aligned} 944 \quad \mathbb{E}_{\mathcal{S}} \left[ \frac{1}{\sigma_{\mathcal{S}}} \right] &\approx \frac{1}{\mathbb{E}_{\mathcal{S}}[\sigma_{\mathcal{S}}]} + \frac{1}{2} \text{Var}_{\mathcal{S}}(\sigma_{\mathcal{S}}) \cdot f''(\mathbb{E}_{\mathcal{S}}[\sigma_{\mathcal{S}}]) \\ 945 \\ 946 &= \frac{1}{\mathbb{E}_{\mathcal{S}}[\sigma_{\mathcal{S}}]} + \frac{1}{2} \text{Var}_{\mathcal{S}}(\sigma_{\mathcal{S}}) \cdot \frac{2}{(\mathbb{E}_{\mathcal{S}}[\sigma_{\mathcal{S}}])^3} \\ 947 \\ 948 &= \frac{1}{\mathbb{E}_{\mathcal{S}}[\sigma_{\mathcal{S}}]} + \frac{\text{Var}_{\mathcal{S}}(\sigma_{\mathcal{S}})}{(\mathbb{E}_{\mathcal{S}}[\sigma_{\mathcal{S}}])^3} \end{aligned} \quad (14)$$

950 This result shows how the expectation of the inverse is related to the inverse of the expectation  
951 plus a correction term that is directly proportional to the variance of the random variable.

- 952 • **Derivation for  $\mathbb{E}_{\mathcal{S}} \left[ \frac{\mu_{\mathcal{S}}}{\sigma_{\mathcal{S}}} \right]$ :** We apply the first-order multivariate Taylor expansion on function  
953  $f(\mu_{\mathcal{S}}, \sigma_{\mathcal{S}}) = \mu_{\mathcal{S}}/\sigma_{\mathcal{S}}$ , giving:

$$\begin{aligned} 955 \quad f(\mu_{\mathcal{S}}, \sigma_{\mathcal{S}}) &\approx f(\mathbb{E}_{\mathcal{S}}[\mu_{\mathcal{S}}], \mathbb{E}_{\mathcal{S}}[\sigma_{\mathcal{S}}]) + (\mu_{\mathcal{S}} - \mathbb{E}_{\mathcal{S}}[\mu_{\mathcal{S}}]) f'_{\mu_{\mathcal{S}}}(\mathbb{E}_{\mathcal{S}}[\mu_{\mathcal{S}}], \mathbb{E}_{\mathcal{S}}[\sigma_{\mathcal{S}}]) \\ 956 \\ 957 &\quad + (\sigma_{\mathcal{S}} - \mathbb{E}_{\mathcal{S}}[\sigma_{\mathcal{S}}]) f'_{\sigma_{\mathcal{S}}}(\mathbb{E}_{\mathcal{S}}[\mu_{\mathcal{S}}], \mathbb{E}_{\mathcal{S}}[\sigma_{\mathcal{S}}]) \end{aligned} \quad (15)$$

958 Taking the expectation of both sides, the linear terms vanish since  $\mathbb{E}_{\mathcal{S}}[\mu_{\mathcal{S}} - \mathbb{E}_{\mathcal{S}}[\mu_{\mathcal{S}}]] = 0$  and  
959  $\mathbb{E}_{\mathcal{S}}[\sigma_{\mathcal{S}} - \mathbb{E}_{\mathcal{S}}[\sigma_{\mathcal{S}}]] = 0$ , and  $\mathbb{E}_{\mathcal{S}}[\mu_{\mathcal{S}}] = \mu_{\mathcal{O}}$  (given by the Law of Total Expectation), we get the  
960 approximation for the expectation:

$$961 \quad \mathbb{E}_{\mathcal{S}} \left[ \frac{\mu_{\mathcal{S}}}{\sigma_{\mathcal{S}}} \right] \approx \frac{\mathbb{E}_{\mathcal{S}}[\mu_{\mathcal{S}}]}{\mathbb{E}_{\mathcal{S}}[\sigma_{\mathcal{S}}]} = \frac{\mu_{\mathcal{O}}}{\mathbb{E}_{\mathcal{S}}[\sigma_{\mathcal{S}}]} \quad (16)$$

963 Now, we substitute the expressions of the two expectation terms in Eq. (14) and Eq. (16) back into  
964 the expanded difference formula in Eq. (12). Plugging in the results from the derivations above and  
965 grouping the terms containing  $r_{i,t}$  and  $\mu_{\mathcal{O}}$ , we have:

$$\begin{aligned} 966 \quad \Delta \hat{A}_{i,t} &\approx r_{i,t} \left( \frac{1}{\mathbb{E}_{\mathcal{S}}[\sigma_{\mathcal{S}}]} + \frac{\text{Var}_{\mathcal{S}}(\sigma_{\mathcal{S}})}{(\mathbb{E}_{\mathcal{S}}[\sigma_{\mathcal{S}}])^3} \right) - \left( \frac{\mu_{\mathcal{O}}}{\mathbb{E}_{\mathcal{S}}[\sigma_{\mathcal{S}}]} \right) - \frac{r_{i,t}}{\sigma_{\mathcal{O}}} + \frac{\mu_{\mathcal{O}}}{\sigma_{\mathcal{O}}} \\ 967 \\ 968 &= r_{i,t} \left( \frac{1}{\mathbb{E}_{\mathcal{S}}[\sigma_{\mathcal{S}}]} - \frac{1}{\sigma_{\mathcal{O}}} \right) - \mu_{\mathcal{O}} \left( \frac{1}{\mathbb{E}_{\mathcal{S}}[\sigma_{\mathcal{S}}]} - \frac{1}{\sigma_{\mathcal{O}}} \right) + r_{i,t} \frac{\text{Var}_{\mathcal{S}}(\sigma_{\mathcal{S}})}{(\mathbb{E}_{\mathcal{S}}[\sigma_{\mathcal{S}}])^3} \\ 969 \\ 970 &= (r_{i,t} - \mu_{\mathcal{O}}) \left( \frac{1}{\mathbb{E}_{\mathcal{S}}[\sigma_{\mathcal{S}}]} - \frac{1}{\sigma_{\mathcal{O}}} \right) + r_{i,t} \frac{\text{Var}_{\mathcal{S}}(\sigma_{\mathcal{S}})}{(\mathbb{E}_{\mathcal{S}}[\sigma_{\mathcal{S}}])^3} \end{aligned} \quad (17)$$

To analyze the terms above, we need to find an expression for the remaining expectation term  $\mathbb{E}_S[\sigma_S]$ . Next, we continue to derive it:

- **Derivation for  $\mathbb{E}_S[\sigma_S]$ :** First, we establish a relationship between the expectation of the squared local standard deviation,  $\mathbb{E}_S[\sigma_S^2]$ , and the square of its expectation,  $(\mathbb{E}_S[\sigma_S])^2$ . This identity is derived directly from the fundamental definition of variance:

$$\begin{aligned}
 \text{Var}(\sigma_S) &= \mathbb{E}_S[(\sigma_S - \mathbb{E}_S[\sigma_S])^2] \\
 &= \mathbb{E}_S[\sigma_S^2 - 2\sigma_S\mathbb{E}_S[\sigma_S] + (\mathbb{E}_S[\sigma_S])^2] \\
 &= \mathbb{E}_S[\sigma_S^2] - \mathbb{E}_S[2\sigma_S\mathbb{E}_S[\sigma_S]] + \mathbb{E}_S[(\mathbb{E}_S[\sigma_S])^2] \\
 &= \mathbb{E}_S[\sigma_S^2] - 2\mathbb{E}_S[\sigma_S]\mathbb{E}[\sigma_S] + (\mathbb{E}_S[\sigma_S])^2 \\
 &= \mathbb{E}_S[\sigma_S^2] - (\mathbb{E}_S[\sigma_S])^2
 \end{aligned} \tag{18}$$

Rearranging this identity gives the desired relationship:

$$(\mathbb{E}_S[\sigma_S])^2 = \mathbb{E}_S[\sigma_S^2] - \text{Var}(\sigma_S) \tag{19}$$

Then, according the *Law of Total Variance*, we have the expression for  $\mathbb{E}_S[\sigma_S^2]$ :

$$\mathbb{E}_S[\sigma_S^2] = \sigma_{\mathcal{O}}^2 - \text{Var}_S(\mu_S) \tag{20}$$

By equating Eq. (19) and Eq. (20), we can solve for  $(\mathbb{E}_S[\sigma_S])^2$ :

$$(\mathbb{E}_S[\sigma_S])^2 = \sigma_{\mathcal{O}}^2 - \text{Var}_S(\mu_S) - \text{Var}_S(\sigma_S) \tag{21}$$

Taking the square root gives us the desired expression for the average local standard deviation:

$$\mathbb{E}_S[\sigma_S] = \sqrt{\sigma_{\mathcal{O}}^2 - \text{Var}_S(\mu_S) - \text{Var}_S(\sigma_S)} \tag{22}$$

Substituting this back into Eq. (17) gives its exact expression, forming the basis of our analysis in the following subsection.

$$\Delta \hat{A}_{i,t} \approx \underbrace{(r_{i,t} - \mu_{\mathcal{O}})}_{\text{global-relative reward}} \underbrace{\left( \frac{1}{\mathbb{E}_S[\sigma_S]} - \frac{1}{\sigma_{\mathcal{O}}} \right)}_{\substack{\text{Correction Term 1} \\ \text{scaling factor 1}}} + \underbrace{\frac{r_{i,t}}{\text{Var}_S(\sigma_S)}}_{\substack{\text{reward} \\ \text{scaling factor 2}}} \underbrace{\frac{\text{Var}_S(\sigma_S)}{(\mathbb{E}_S[\sigma_S])^3}}_{\text{Correction Term 2}}, \tag{23}$$

where  $\mathbb{E}_S[\sigma_S] = \sqrt{\sigma_{\mathcal{O}}^2 - \text{Var}_S(\mu_S) - \text{Var}_S(\sigma_S)}$

## D.2 ANALYSIS ON THE PROPERTIES OF “SCALING FACTOR 1” IN EQ. (23)

The “scaling factor 1” in Eq. (23) has two crucial properties:

- It is non-negative.
- Its magnitude increases with the heterogeneity of reward distribution.

Below, we prove these properties in Proof 1 and Proof 2, respectively.

**Proof 1: The “scaling factor 1”  $\left(\mathbb{E}_S\left[\frac{1}{\sigma_S}\right] - \frac{1}{\sigma_{\mathcal{O}}}\right)$  is non-negative**

This property can be formally proven using the expression for  $\mathbb{E}_S[\sigma_S]$  derived in Eq. (23), i.e.,  $\mathbb{E}_S[\sigma_S] = \sqrt{\sigma_{\mathcal{O}}^2 - \text{Var}_S(\mu_S) - \text{Var}_S(\sigma_S)}$ . Since both variance terms,  $\text{Var}_S(\mu_S)$  and  $\text{Var}_S(\sigma_S)$ , are non-negative by definition, the term inside the square root is always less than or equal to  $\sigma_{\mathcal{O}}^2$ . This leads to the inequality  $\mathbb{E}_S[\sigma_S] \leq \sigma_{\mathcal{O}}$ . Taking the reciprocal of both sides reverses the inequality, which in turn proves that:  $\mathbb{E}_S\left[\frac{1}{\sigma_S}\right] - \frac{1}{\sigma_{\mathcal{O}}} \geq 0$ .

**Proof 2: The magnitude of “scaling factor 1”  $\left(\mathbb{E}_S\left[\frac{1}{\sigma_S}\right] - \frac{1}{\sigma_{\mathcal{O}}}\right)$  increases with the heterogeneity of reward distribution.**

The concept of *reward heterogeneity* refers to the degree of non-uniformity in the reward distribution within group  $\mathcal{O}$ . High heterogeneity implies an uneven spread of rewards, potentially characterized by significant outliers or the emergence of distinct clusters, such as a “high-reward” subgroup or a “low-reward” subgroup. Reward heterogeneity can be quantitatively captured by two key terms:  $\text{Var}_{\mathcal{S}}(\mu_{\mathcal{S}})$  and  $\text{Var}_{\mathcal{S}}(\sigma_{\mathcal{S}})$ .

- First,  $\text{Var}_{\mathcal{S}}(\mu_{\mathcal{S}})$ , the variance of subgroup means, serves as the primary and most direct measure of heterogeneity. As revealed by the Law of Total Variance ( $\sigma_{\mathcal{O}}^2 = \mathbb{E}_{\mathcal{S}}[\sigma_{\mathcal{S}}^2] + \text{Var}_{\mathcal{S}}(\mu_{\mathcal{S}})$ ),  $\text{Var}_{\mathcal{S}}(\mu_{\mathcal{S}})$  accounts for the portion of global variance attributable to differences between subgroup averages. It thus reflects how much the “center of gravity” of rewards shifts across local views. A large  $\text{Var}_{\mathcal{S}}(\mu_{\mathcal{S}})$  is a strong indicator of clustered reward structures.
- Second,  $\text{Var}_{\mathcal{S}}(\sigma_{\mathcal{S}})$ , the variance of subgroup standard deviations, acts as a secondary but informative signature of heterogeneity. In heterogeneous populations, subgroups exhibit highly variable internal dispersion, for instance, a homogeneous high-reward subgroup will have low  $\sigma_{\mathcal{S}}$ , while a mixed subgroup may display high variability. The fluctuation in these local standard deviations across subgroups is precisely what  $\text{Var}_{\mathcal{S}}(\sigma_{\mathcal{S}})$  measures.

In conclusion, as reward heterogeneity increases,  $\text{Var}_{\mathcal{S}}(\mu_{\mathcal{S}})$  and  $\text{Var}_{\mathcal{S}}(\sigma_{\mathcal{S}})$  grow, widening the gap between  $\sigma_{\mathcal{O}}^2$  and  $\mathbb{E}_{\mathcal{S}}[\sigma_{\mathcal{S}}^2]$ , and thereby increasing the magnitude of  $(1/\mathbb{E}_{\mathcal{S}}[\sigma_{\mathcal{S}}] - 1/\sigma_{\mathcal{O}})$ . This establishes a clear link between structural heterogeneity in rewards and the magnitude of the second factor  $(1/\mathbb{E}_{\mathcal{S}}[\sigma_{\mathcal{S}}] - 1/\sigma_{\mathcal{O}})$ .

### D.3 ANALYSIS ON THE PROPERTIES OF “SCALING FACTOR 2” IN EQ. (23)

The “scaling factor 2” in Eq. (23) has two crucial properties:

- It is non-negative.
- Its magnitude increases with the heterogeneity of reward distribution.

Below, we prove these properties in Proof 3 and Proof 4, respectively.

#### Proof 3: The “scaling factor 2” $\left(\frac{\text{Var}_{\mathcal{S}}(\sigma_{\mathcal{S}})}{(\mathbb{E}_{\mathcal{S}}[\sigma_{\mathcal{S}}])^3}\right)$ is non-negative

First, the numerator term  $\text{Var}_{\mathcal{S}}(\sigma_{\mathcal{S}})$  represents the variance of the random variable  $\sigma_{\mathcal{S}}$ . By the definition of variance,  $\text{Var}_{\mathcal{S}}(\sigma_{\mathcal{S}})$  is always non-negative, *i.e.*,  $\text{Var}_{\mathcal{S}}(\sigma_{\mathcal{S}}) \geq 0$ .

Second, for the denominator term  $(\mathbb{E}_{\mathcal{S}}[\sigma_{\mathcal{S}}])^3$ , to determine the sign, we first analyze its base,  $\mathbb{E}_{\mathcal{S}}[\sigma_{\mathcal{S}}]$ . Taking the expression from Eq. (23), we can know that  $\mathbb{E}_{\mathcal{S}}[\sigma_{\mathcal{S}}] = \sqrt{\sigma_{\mathcal{O}}^2 - \text{Var}_{\mathcal{S}}(\mu_{\mathcal{S}}) - \text{Var}_{\mathcal{S}}(\sigma_{\mathcal{S}})} \geq 0$ . Thus, its cube,  $(\mathbb{E}_{\mathcal{S}}[\sigma_{\mathcal{S}}])^3$ , must also be non-negative, *i.e.*,  $(\mathbb{E}_{\mathcal{S}}[\sigma_{\mathcal{S}}])^3 \geq 0$ .

In conclusion, since both the numerator  $\text{Var}_{\mathcal{S}}(\sigma_{\mathcal{S}})$  and denominator  $(\mathbb{E}_{\mathcal{S}}[\sigma_{\mathcal{S}}])^3$  are non-negative, their ratio is guaranteed to be non-negative. We have proven that  $\frac{\text{Var}_{\mathcal{S}}(\sigma_{\mathcal{S}})}{(\mathbb{E}_{\mathcal{S}}[\sigma_{\mathcal{S}}])^3} \geq 0$ .

#### Proof 4: The magnitude of “scaling factor 2” $\left(\frac{\text{Var}_{\mathcal{S}}(\sigma_{\mathcal{S}})}{(\mathbb{E}_{\mathcal{S}}[\sigma_{\mathcal{S}}])^3}\right)$ increases with the heterogeneity of reward distribution.

As discussed in Proof 2, the reward heterogeneity can be quantitatively measured by terms such as  $\text{Var}_{\mathcal{S}}(\mu_{\mathcal{S}})$  and  $\text{Var}_{\mathcal{S}}(\sigma_{\mathcal{S}})$ . As reward heterogeneity increases,  $\text{Var}_{\mathcal{S}}(\mu_{\mathcal{S}})$  and  $\text{Var}_{\mathcal{S}}(\sigma_{\mathcal{S}})$  also grows.

Recalling the expression of  $\mathbb{E}_{\mathcal{S}}[\sigma_{\mathcal{S}}]$  derived in Eq. (23), we have  $\mathbb{E}_{\mathcal{S}}[\sigma_{\mathcal{S}}] = \sqrt{\sigma_{\mathcal{O}}^2 - \text{Var}_{\mathcal{S}}(\mu_{\mathcal{S}}) - \text{Var}_{\mathcal{S}}(\sigma_{\mathcal{S}})}$ . As reward heterogeneity increases,  $\text{Var}_{\mathcal{S}}(\mu_{\mathcal{S}})$  and  $\text{Var}_{\mathcal{S}}(\sigma_{\mathcal{S}})$  also grows, thereby leading to the corresponding decrease of  $\mathbb{E}_{\mathcal{S}}[\sigma_{\mathcal{S}}]$ . Since  $\mathbb{E}_{\mathcal{S}}[\sigma_{\mathcal{S}}] > 0$ , its cube,  $(\mathbb{E}_{\mathcal{S}}[\sigma_{\mathcal{S}}])^3$  decreases as well. Thus, as the reward heterogeneity increases, the numerator  $\text{Var}_{\mathcal{S}}(\sigma_{\mathcal{S}})$  increases, while the denominator  $(\mathbb{E}_{\mathcal{S}}[\sigma_{\mathcal{S}}])^3$  declines, thereby resulting in a larger  $\frac{\text{Var}_{\mathcal{S}}(\sigma_{\mathcal{S}})}{(\mathbb{E}_{\mathcal{S}}[\sigma_{\mathcal{S}}])^3}$ . This formally proves that the magnitude of “scaling factor 2”,  $\frac{\text{Var}_{\mathcal{S}}(\sigma_{\mathcal{S}})}{(\mathbb{E}_{\mathcal{S}}[\sigma_{\mathcal{S}}])^3}$ , is a increasing function of the heterogeneity of the reward distribution.

1080 **E APPROXIMATION GUARANTEE FOR THE GREEDY ALGORITHM**  
 1081

1082 In this section, we provide a theoretical justification for the greedy algorithm used for the *Diverse*  
 1083 *Group Sampling* procedure (Sec. 3.2.2). We first formalize the problem and then provide a proof  
 1084 sketch for the approximation guarantee in the context of submodular function maximization, to  
 1085 which our diversity objective is closely related.

1087 **E.1 PROBLEM FORMULATION**  
 1088

1089 As stated in the main text, for a given scale  $\tau$ , we aim to select a subset of comparison groups  $\mathbb{C}_\tau$  of  
 1090 size  $K$  from an exhaustive set  $\mathbb{S}_\tau$  that solves the Maximum Diversity problem:

$$1091 \quad \mathbb{C}_\tau^* = \arg \max_{\mathbb{C}_\tau \subset \mathbb{S}_\tau, |\mathbb{C}_\tau|=K} \sum_{\mathcal{S}_a, \mathcal{S}_b \in \mathbb{C}_\tau, a \neq b} d_J(\mathcal{S}_a, \mathcal{S}_b). \quad (24)$$

1094 Let  $f(\mathbb{C}_\tau) = \sum_{\mathcal{S}_a, \mathcal{S}_b \in \mathbb{C}_\tau, a \neq b} d_J(\mathcal{S}_a, \mathcal{S}_b)$  be our objective function. As this problem is NP-hard, we  
 1095 employ a greedy algorithm that iteratively constructs the set. Let  $\mathbb{C}_k$  be the set of  $k$  groups selected  
 1096 after  $k$  iterations. In step  $k+1$ , the algorithm selects the group  $\mathcal{S}_{k+1}$  that provides the maximum  
 1097 marginal gain:

$$1098 \quad \mathcal{S}_{k+1} = \arg \max_{\mathcal{S} \in \mathbb{S}_\tau \setminus \mathbb{C}_k} f(\mathbb{C}_k \cup \{\mathcal{S}\}) - f(\mathbb{C}_k). \quad (25)$$

1100 Note that the selection criterion in the main text,  $\sum_{\mathcal{S}' \in \mathbb{C}_k} d_J(\mathcal{S}, \mathcal{S}')$ , is precisely this marginal gain.

1102 **E.2 SUBMODULARITY AND APPROXIMATION GUARANTEES**  
 1103

1104 The effectiveness of this greedy strategy is best understood through the lens of *submodularity*. A set  
 1105 function  $f : 2^V \rightarrow \mathbb{R}$  is submodular if for any two sets  $A \subseteq B \subset V$  and any element  $x \in V \setminus B$ , it  
 1106 satisfies the “diminishing returns” property:

$$1107 \quad f(A \cup \{x\}) - f(A) \geq f(B \cup \{x\}) - f(B). \quad (26)$$

1109 In words, the marginal gain of adding an element  $x$  to a small set  $A$  is greater than or equal to the  
 1110 marginal gain of adding the same element to a larger superset  $B$ .

1111 While our Max-Sum Diversity objective  $f(\mathbb{C}_\tau)$  is not strictly submodular, it is closely related to this  
 1112 class of functions, and the greedy algorithm is the principled approach for both. For the general class  
 1113 of non-negative, monotone submodular functions, a celebrated result provides a strong performance  
 1114 guarantee for the greedy algorithm.

1115 The theorem in (Nemhauser et al., 1978) states that: for the problem of maximizing a non-negative,  
 1116 monotone submodular function  $f$  subject to a cardinality constraint  $|\mathbb{C}| \leq K$ , the greedy algorithm  
 1117 produces a solution  $\mathbb{C}_G$  such that:  $f(\mathbb{C}_G) \geq (1 - \frac{1}{e}) f(\mathbb{C}^*)$ , where  $\mathbb{C}^*$  is the true optimal solution  
 1118 and  $e$  is the base of the natural logarithm. This guarantees that the greedy solution is within  
 1119 approximately 63.2% of the optimal solution.

1121 **E.3 PROOF SKETCH FOR THE SUBMODULAR CASE**  
 1122

1123 We provide a sketch of the classic proof for the  $(1 - 1/e)$  approximation guarantee to illustrate  
 1124 the principle. Let  $\mathbb{C}_k$  be the greedy set of size  $k$  and  $\mathbb{C}^*$  be the optimal set of size  $K$ . Let  $g_k =$   
 1125  $f(\mathbb{C}_k) - f(\mathbb{C}_{k-1})$  be the marginal gain at step  $k$ .

1126 The key insight is to bound the difference between the optimal value and the current greedy value.  
 1127 By monotonicity and submodularity, we can write:

$$1128 \quad f(\mathbb{C}^*) \leq f(\mathbb{C}_k \cup \mathbb{C}^*) \quad (27)$$

$$1130 \quad = f(\mathbb{C}_k) + \sum_{x \in \mathbb{C}^* \setminus \mathbb{C}_k} (f(\mathbb{C}_k \cup \{x\}) - f(\mathbb{C}_k)) \quad (28)$$

$$1132 \quad \leq f(\mathbb{C}_k) + \sum_{x \in \mathbb{C}^* \setminus \mathbb{C}_k} (f(\mathbb{C}_{k-1} \cup \{x\}) - f(\mathbb{C}_{k-1})) \quad (\text{by submodularity}) \quad (29)$$

1134 The greedy choice at step  $k$  maximizes the marginal gain, so its gain  $g_k$  is at least the average gain  
 1135 of the elements in  $\mathbb{C}^* \setminus \mathbb{C}_{k-1}$ . This leads to the inequality:  
 1136

$$f(\mathbb{C}^*) - f(\mathbb{C}_{k-1}) \leq K \cdot g_k = K \cdot (f(\mathbb{C}_k) - f(\mathbb{C}_{k-1})) \quad (30)$$

1138 Rearranging this recurrence relation over  $k$  from 1 to  $K$  and using the fact that  $(1 - 1/K)^K \approx 1/e$   
 1139 leads to the final bound  $f(\mathbb{C}_K) \geq (1 - 1/e)f(\mathbb{C}^*)$ .  
 1140

1141 This proof sketch demonstrates how the diminishing returns property of submodularity allows the  
 1142 greedy algorithm to provide a constant-factor approximation of the true optimum. While our specific  
 1143 objective requires a more tailored analysis, this classic result provides the theoretical foundation  
 1144 for why the greedy approach is a fast, principled, and provably effective choice for our diversity  
 1145 maximization task.  
 1146  
 1147  
 1148  
 1149  
 1150  
 1151  
 1152  
 1153  
 1154  
 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