

---

# HARD EXAMPLES ARE ALL YOU NEED: MAXIMIZING GRPO POST-TRAINING UNDER ANNOTATION BUDGETS

**Anonymous authors**

Paper under double-blind review

## ABSTRACT

Collecting high-quality training examples for language model fine-tuning is expensive, with practical budgets limiting the amount of data that can be procured. We investigate whether example difficulty affects GRPO training effectiveness by comparing selection strategies (easy, medium, hard, random) across multiple models and reasoning tasks. Training on the hardest 10% of examples (those where the base model fails most often) yields dramatic performance gains up to 47%, while easy examples produce minimal improvements of 3-15%. This occurs because GRPO requires outcome variance to generate learning signals; hard examples maintain mixed success/failure outcomes throughout training while easy examples quickly converge to consistent success, eliminating learning opportunities. Moreover, models trained on hard examples show superior out-of-distribution generalization, with only hard-trained models achieving meaningful gains on the AIME2025 benchmark. Our findings provide clear guidance: when budget-constrained, prioritize collecting and annotating examples where your base model struggles, as these drive nearly all learning value in GRPO fine-tuning.

## 1 INTRODUCTION

Large language model (LLM) alignment and post-training are fundamentally constrained by the expense of acquiring high-quality supervision data, making the selection of training examples a critical factor in achieving optimal model performance Ouyang et al. (2022).<sup>1</sup> While recent advances in reinforcement learning from human feedback (RLHF) and related techniques have demonstrated remarkable improvements in model capabilities, practitioners face a practical challenge: given limited resources for data annotation and curation, which examples should be prioritized to maximize post-training performance?

In this work, we focus on Group Relative Policy Optimization (GRPO), a PPO-style algorithm that replaces a learned value function with group-normalized advantages, reducing memory and relying on within-group reward variance for learning signals. We address a specific instantiation of this budget-aware selection problem: given a fixed budget to select and train on only a fraction of available prompts, how should we choose this subset to maximize the effectiveness of Group Relative Policy Optimization (GRPO) fine-tuning? Specifically, we investigate whether selecting examples based on their difficulty (as measured by the base model’s success rate across multiple sampling attempts) leads to systematic differences in final model performance on held-out test sets.

Our key research questions are:

1. Should practitioners prioritize examples that are hard (where the base model frequently fails), easy (where it frequently succeeds), medium, or a selection of random difficulty examples?
2. Does the optimal selection strategy vary across model scales and families?
3. What mechanisms explain the differential effectiveness of difficulty-based selection strategies under GRPO’s learning dynamics?

---

<sup>1</sup>Code available here

- 
- 054 4. How do different difficulty-based selection strategies impact out-of-distribution generaliza-  
055 tion to substantially harder problem sets?
  - 056 5. Can we achieve similar performance gains by training only on examples the base model gets  
057 wrong, avoiding the circular dependency of knowing which examples will improve?  
058

059 Our findings reveal a striking pattern: **hard examples are all you need**. Training on the hardest  
060 10% of examples yields performance gains up to 47%, while easy examples produce minimal  
061 improvements of 3-15%. This 30+ percentage point gap transforms marginally effective fine-tuning  
062 into highly successful model improvement.

063 To answer these questions, we make the following contributions:  
064

065 **Budget-aware evaluation protocol.** We develop a systematic framework for comparing difficulty-  
066 conditioned training subsets under GRPO, using multi-sample base-model probing to robustly  
067 estimate example difficulty Wang et al. (2022). This protocol enables fair comparison across selection  
068 strategies while controlling for computational budget and other confounding factors.  
069

070 **Comprehensive experimental evaluation.** We conduct extensive experiments on GSM8K grade-  
071 school math problems and BIG-Bench Hard’s Tracking Shuffled Objects task across models (Qwen3-  
072 4B, Qwen3-14B, Phi-4, and Llama3.1-8B) Dubey et al. (2024); Yang et al. (2025); Abdin et al.  
073 (2024). Our results consistently show that training on the hardest 10% of examples yields superior  
074 test performance on reasoning tasks.

075 **Mechanistic understanding of selection effects.** We analyze why hard examples prove most  
076 effective under GRPO: the algorithm requires within-group outcome variance to generate learning  
077 signals. When all samples in a group produce identical rewards, the advantages become zero and  
078 learning stops. Our analysis shows that hard examples maintain mixed outcomes longer than easy  
079 examples, which quickly converge to deterministic success Shao et al. (2024).  
080

081 **Out-of-distribution generalization analysis.** We evaluate models trained on different difficulty-  
082 based subsets against the AIME2025-I benchmark, a substantially harder test set than the in-  
083 distribution GSM8K OpenCompass (2025); Balunović et al. (2025). We find that training on the  
084 hardest examples not only improves in-distribution performance, but is also the only strategy that  
085 yields meaningful gains on OOD data.  
086

087 **Base wrong vs base right analysis.** We validate our “hard examples” principle from a complementary  
088 perspective by categorizing examples based on base model performance. This analysis reveals that  
089 nearly all learning value comes from examples the base model initially gets wrong, with training on  
090 “base wrong” examples achieving up to 23.5% relative improvements over “base right” examples.  
091 This convergent evidence that both low pass@k examples and base-wrong examples drive learning  
092 confirms that GRPO fundamentally benefits from training at the frontier of model capabilities.

093 Our findings have immediate practical implications: when using GRPO for reasoning task fine-tuning  
094 under budget constraints, practitioners should prioritize collecting and annotating examples where  
095 the base model struggles. We validate this principle from two complementary angles: examples with  
096 low pass@k success rates and examples the base model gets wrong both drive superior learning  
097 outcomes. This convergence demonstrates that “hard examples” identified through either metric  
098 capture the same underlying phenomenon: problems at the frontier of model capabilities that provide  
099 rich learning signals. This hard-example focus can yield performance improvements *exceeding 30*  
100 *percentage points* over easy selection on some configurations, transforming marginally effective  
101 fine-tuning into highly successful model improvement.

102 Our study isolates *offline* difficulty-based selection under GRPO. It complements *online*  
103 difficulty-targeted selection with rollout replay that prioritizes moderate difficulty Sun et al. (2025),  
104 advances in process-reward RL (PRIME) Cui et al. (2025), and analyses of GRPO optimization biases  
105 such as Dr. GRPO Liu et al. (2025). We focus on outcome-reward GRPO; extending our protocol to  
106 dense process rewards or debiased GRPO is left to future work.

107 The remainder of this paper is organized as follows. Section 2 presents our experimental protocol for  
difficulty estimation and subset selection. Section 3 reports our main experimental results across tasks

and model scales. Section 4 provides detailed analysis of learning dynamics and model behavior. Section 5 concludes with practical recommendations and future directions. Related work can be found in the appendix.

## 2 EXPERIMENTAL PROTOCOL

This section describes our experimental framework for evaluating the impact of example difficulty on GRPO fine-tuning under budget constraints. We detail our difficulty estimation procedure, subset selection policies, training methodology, and evaluation metrics.

### 2.1 GOAL

Our primary research goal is to determine which subset selection strategy maximizes post-GRPO test accuracy when constrained to using only  $p = 10\%$  of the available training pool. Specifically, given a pool of unlabeled prompts  $\mathcal{X}$ , we select subsets  $S \subset \mathcal{X}$  with  $|S| = \lfloor p|\mathcal{X}| \rfloor$  according to different difficulty-based policies, train models using GRPO on each subset, and compare their performance on held-out test sets. This setup mirrors practical scenarios where annotation budgets limit the number of examples that can be used for fine-tuning.

### 2.2 DIFFICULTY ESTIMATION VIA MULTI-SAMPLE PROBING

To robustly estimate example difficulty, we employ multi-sample evaluation using the base model before any fine-tuning. For each prompt  $x$  in the training pool, we sample  $K$  independent completions from the base model  $\pi_{\text{base}}$  using temperature  $\tau = 1.0$  and chain-of-thought prompting where applicable Wang et al. (2022). For each prompt  $x$ , we compute:

$$\hat{p}(x) = \frac{1}{K} \sum_{i=1}^K \mathbf{1}[\text{completion}_i \text{ is correct}] \quad (1)$$

where  $\hat{p}(x)$  represents the empirical success rate.

#### Task-specific correctness criteria:

- **GSM8K:** Exact match of the final numerical answer after extracting from the generated solution Cobbe et al. (2021)
- **Tracking Shuffled Objects:** Correct identification of all object positions after the described swaps Suzgun et al. (2023)

We use  $K = 10$  samples for Tracking Shuffle Objects, and  $K = 5$  for GSM8K due to its larger size and associated computational costs. This multi-sample approach provides stable difficulty estimates while remaining computationally tractable.

### 2.3 SUBSET SELECTION POLICIES

Given difficulty estimates for all prompts, we implement four selection policies, each choosing exactly 10% of the training pool:

1. **HARD(-est):** Select prompts with the lowest success rates:

$$S_{\text{hard}} = \arg \min_{S: |S| = \lfloor p|\mathcal{X}| \rfloor} \sum_{x \in S} \hat{p}(x)$$

2. **EASY(-est):** Select prompts with the highest success rates:

$$S_{\text{easy}} = \arg \max_{S: |S| = \lfloor p|\mathcal{X}| \rfloor} \sum_{x \in S} \hat{p}(x)$$

3. **MEDIUM:** Select prompts nearest to the median difficulty, specifically those in the interquartile range around the median  $\hat{p}(x)$  value
4. **RANDOM:** Uniform random sample without replacement from the full pool

Figure 1 illustrates the complete pipeline from difficulty estimation through subset selection to GRPO training.

### Budget-Aware GRPO Pipeline

1. **Difficulty Estimation:** Sample  $K$  completions per prompt from base model  $\rightarrow$  Compute success rate  $\hat{p}(x)$
2. **Subset Selection:** Apply policy (HARD, EASY, MEDIUM, RANDOM)  $\rightarrow$  Select 10% subset
3. **GRPO Training:** Fine-tune model on selected subset using group advantages
4. **Evaluation:** Test on held-out benchmark

Figure 1: Schematic overview of our experimental protocol

## 2.4 TRAINING PROCEDURE

We employ Group Relative Policy Optimization (GRPO) following the formulation introduced in DeepSeekMath Shao et al. (2024). More details on the specifics of the GRPO algorithm and hyperparameters used can be found in the appendix.

Our experiments span two benchmark datasets that probe different aspects of model capabilities. GSM8K provides 7,473 grade-school math problems requiring multi-step reasoning and arithmetic Cobbe et al. (2021). The Tracking 7 Shuffled Objects task from BIG-Bench Hard contains 250 problems testing the ability to maintain state through sequences of object swaps Suzgun et al. (2023), which we randomly split 50%/50% into train/test. We hold the number of prompts, rollouts per prompt, group size, KL schedule, max length, and total RL update steps constant across policies. This avoids claims being confounded by extra tokens/updates.

We evaluate three models: Qwen3-4B with 4 billion parameters representing a smaller but efficient architecture, Qwen3-14B with 14 billion parameters as a medium-scale model, and Phi-4 with 14 billion parameters, representing a different model family. All models are initialized from instruction-tuned checkpoints, ensuring they have basic instruction-following capabilities before GRPO fine-tuning begins.

## 3 EXPERIMENTS AND RESULTS

We present our main experimental findings comparing difficulty-based subset selection strategies for GRPO fine-tuning. Our results consistently demonstrate that training on the hardest examples yields superior performance on reasoning tasks.

### 3.1 HARD EXAMPLES DRIVE REASONING IMPROVEMENTS

Figure 2 and Table 3.1, summarize the absolute accuracy change (in percentage points) from the baseline after GRPO training, across selection policies, datasets, and models.

Dataset	Model	EASY	MEDIUM	HARD	RANDOM
GSM8K	Qwen3-4B	3.49	26.69	<b>34.19</b>	<u>29.49</u>
	Qwen3-14B	8.26	24.34	<b>39.42</b>	<u>34.87</u>
	Phi-4	14.94	28.51	<b>37.30</b>	<u>36.16</u>
Shuffled Objects	Qwen3-4B	0.80	-3.20	<b>13.60</b>	-1.60
	Qwen3-14B	3.20	10.12	<b>22.40</b>	<u>12.52</u>
	Phi-4	<u>29.60</u>	28.80	<b>32.80</b>	28.80

Table 1: Absolute performance change (%) after GRPO training with different subset selection policies compared to baseline. Bold indicates best performance, underline indicates second-best.

The results reveal striking patterns in how difficulty-based selection impacts GRPO training effectiveness. On GSM8K, the superiority of hard example selection is unambiguous across all models. The HARD policy achieves remarkable performance improvements of 34.19% for Qwen3-4B, 39.42% for Qwen3-14B, and 37.3% for Phi-4. These gains dwarf those achieved by easy selection, which manages only 3.49%, 8.26%, and 14.94% respectively. The magnitude of these differences (exceeding 30 percentage points) suggests that training on challenging examples fundamentally alters the learning

216  
217  
218  
219  
220  
221  
222  
223  
224  
225  
226  
227  
228  
229  
230  
231  
232  
233  
234  
235  
236  
237  
238  
239  
240  
241  
242  
243  
244  
245  
246  
247  
248  
249  
250  
251  
252  
253  
254  
255  
256  
257  
258  
259  
260  
261  
262  
263  
264  
265  
266  
267  
268  
269

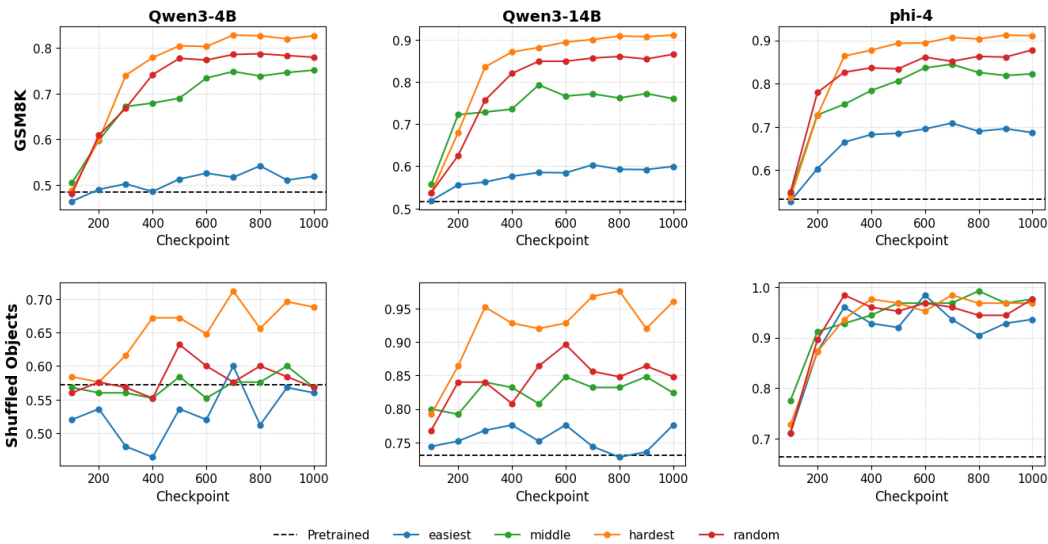


Figure 2: GRPO training dynamics reveal early and persistent advantages of hard example selection. Each subplot shows test accuracy over 1000 training steps for models trained on different difficulty-based subsets (10% of full data each). The hardest subset establishes superiority by step 300 and maintains this advantage.

dynamics under GRPO. We see training on either medium or random difficulties yield significant gains compared to training on easy examples, but not near the gains of training on the hardest subset - affirming the hypothesis that the harder the training subset, the larger the gains we see (Cobbe et al. (2021)).

The Tracking Shuffled Objects task amplifies these patterns. Hard selection yields positive gains of 13.60%, 22.40%, and 32.80% for Qwen3-4B, Qwen3-14B, and Phi-4 respectively. For Qwen3-4B, we see all other subsets produce either zero or negative gains. For Qwen3-14B, improvements show a stepwise pattern: training on the easy subset performs poorly, medium and random subsets provide moderate gains, and the hard subset delivers the largest gains. For Phi-4, training on the hardest subset gives the highest gains, but training on any subset still produces substantial improvements, effectively saturating the benchmark (Suzgun et al. (2023)).

### 3.2 PERFORMANCE DYNAMICS DURING TRAINING

Figure 2 reveals how selection strategies shape learning trajectories from the earliest stages of training. The advantage of hard example selection manifests rapidly and decisively: by checkpoint 300, models trained on the hardest subset have already established clear superiority over all other selection strategies. This early divergence is particularly striking given that it represents less than a third of the total training steps, yet the performance ordering established at this point persists and often amplifies through the remainder of training. The consistency of this pattern across all three models suggests that hard examples provide fundamentally different learning signals that reshape the optimization landscape from the outset.

## 4 ANALYSIS

In this section, we try to better understand why hard examples are most effective for GRPO training in three ways. First, we examine learning dynamics through the lens of new metrics - “learnable percentage”. Second, we inspect performance on an out-of-distribution set to measure generalization capabilities when trained on different subsets. Third, we look at training just on examples where the base model is wrong vs where it is right.

270 4.1 WHY HARD EXAMPLES HELP: MAINTAINING LEARNABLE SAMPLES  
 271  
 272

273 A key insight into GRPO’s preference for hard examples comes from tracking what fraction of  
 274 training examples maintain non-zero outcome variance during training. Recall that GRPO learns  
 275 by contrasting outputs within a group; when all outputs in a group produce identical rewards (aka  
 276 standard deviation of reward is zero), the advantages become zero and no learning occurs.

277 **Learnable percentage.** Let  $T$  be the number of training updates. At update  $t$ , let  $\mathbf{r}_t =$   
 278  $(r_{t,1}, \dots, r_{t,G}) \in \{0, 1\}^G$  denote the vector of outcome rewards for the  $G$  rollouts in the GRPO group  
 279 under  $\pi_t$ . We define the *learnable percentage* as

280  
 281  
 282 
$$\text{Learnable \%} = 100 \cdot \frac{1}{T} \sum_{t=1}^T \mathbf{1}\{\text{StdDev}(\mathbf{r}_t) > 0\},$$
  
 283  
 284  
 285

286 i.e., the fraction of training steps whose within-group reward standard deviation is nonzero (and  
 287 therefore advantage is nonzero).

288 Table 2 shows the % learnable across strategies and datasets. We see that, unsurprisingly, the hardest  
 289 strategies have the highest % learnable examples, meaning training on this data subset gives the most  
 290 opportunity for the model to learn during training. Conversely, we see that the easy strategy has  
 291 very few learnable examples. To test whether this effect depends on the model’s baseline ability, we  
 292 evaluate Llama-3.1-8B-Instruct Dubey et al. (2024), which has much lower initial performance on  
 293 GSM8K. With a weaker baseline, we expect smaller performance gaps between subsets. Table 3  
 294 shows that while the hardest subset still achieves the largest gain, it is closer to the other subsets.  
 295 Notably, all subsets have much higher % learnable values than the corresponding Qwen results (e.g.,  
 296 for the easiest subset, 21% vs. 4% for Qwen3-4B).  
 297  
 298

299

Strategy	Qwen3-4B		Qwen3-14B		Phi-4	
	shuffleobj	gsm8k	shuffleobj	gsm8k	shuffleobj	gsm8k
Easy	9.00	3.70	2.30	7.80	14.20	6.60
Medium	5.70	<u>24.50</u>	7.20	<u>18.90</u>	15.30	<u>23.30</u>
Hard	<b>14.40</b>	<b>34.10</b>	<b>13.50</b>	<b>29.50</b>	<u>19.10</u>	<b>40.20</b>
Random	<u>11.50</u>	19.00	<u>8.00</u>	18.40	<b>19.30</b>	21.50

300  
 301  
 302  
 303  
 304  
 305  
 306 Table 2: % Learnable across models, strategies and datasets  
 307  
 308  
 309  
 310

311

Strategy	Improvement over Base	% Learnable
Easy	36.54	20.90
Medium	39.42	42.00
Hard	<b>40.33</b>	<b>55.20</b>
Random	39.35	39.30

312  
 313  
 314  
 315  
 316  
 317 Table 3: Absolute accuracy improvement and % learnable for Llama-3.1-8B-Instruct on GSM8K  
 318 dataset  
 319  
 320

321 Figure 3 plots the relationship between the % learnable and performance improvement, across all the  
 322 models and strategies. We see a strong positive correlation, confirming that the underlying reason  
 323 the harder subsets give better results is because they maintain more learnable examples throughout  
 training.

324  
325  
326  
327  
328  
329  
330  
331  
332  
333  
334  
335  
336  
337  
338  
339  
340  
341  
342  
343  
344  
345  
346  
347  
348  
349  
350  
351  
352  
353  
354  
355  
356  
357  
358  
359  
360  
361  
362  
363  
364  
365  
366  
367  
368  
369  
370  
371  
372  
373  
374  
375  
376  
377

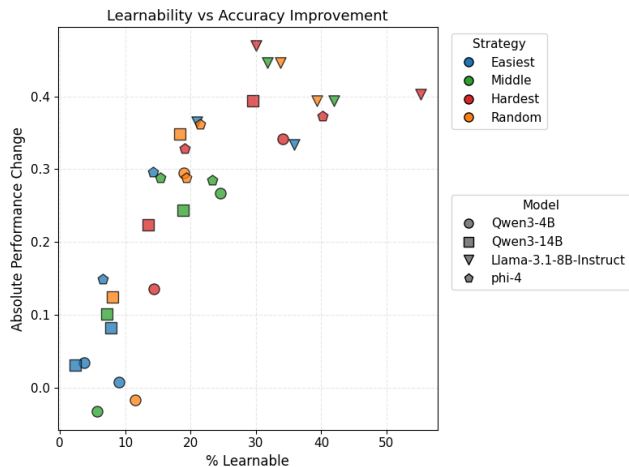


Figure 3: Scatter plot showing the relationship between the percentage of learnable training examples and the resulting absolute improvement in model performance, across strategies and models. Colors indicate the strategy and marker shapes indicate the model. We see a strong positive correlation ( $R^2 = 0.66$ ), indicating that performance improves with more learnable examples.

Our analysis reveals why hard example selection proves dramatically more effective for GRPO training. Hard examples continue providing useful learning signal throughout training, while easy examples quickly become “solved” and stop contributing to improvement. This sustained learning opportunity explains the remarkable performance gaps we observe (often exceeding 30 percentage points between hard and easy selection strategies).

#### 4.2 OUT-OF-DISTRIBUTION PERFORMANCE

Having shown that harder training subsets improve test-set performance, we further evaluate on an out-of-distribution (OOD) test set. Specifically, we take the Qwen3-4B models trained on each subset and evaluate on AIME2025-I, a dataset of 15 questions significantly more difficulty than GSM8K Balunović et al. (2025); OpenCompass (2025). Table 4 presents the results.

Strategy	Pass@8
Base	33.3%
Easy	33.3%
Medium	26.7%
Hard	<b>40.0%</b>
Random	33.3%

Table 4: Pass@8 on AIME-2025 of Qwen3-4B trained on different subsets

We see that the training on the hardest subset is the only one to outperform the base Qwen3-4B model. Surprisingly, training on the medium subset actually does worse than training on easy or random subsets.

We see that training on the hardest subset is the only condition that meaningfully outperforms the base Qwen3-4B model, achieving a relative improvement of +20%. This reinforces the trend observed on the in-distribution test set: exposure to more challenging problems during training generalizes better to harder OOD problems. Interestingly, overall performance on the easy and random subsets remains identical to the base model, with closer inspection showing that these subsets improved performance on certain problems but regressed performance in other areas. In contrast, training on the medium difficulty subset yields a notable drop in performance, with no other gains.

Overall, these results show that the benefits of hard-example training extend beyond the original distribution, improving robustness to unseen, more challenging problems.

### 4.3 BASE WRONG VS BASE RIGHT: DECOMPOSING GRPO VALUE

Having established that hard examples (those with low pass@k success rates) drive the most value in GRPO training, we now validate this principle from a complementary perspective. We investigate whether the superiority of hard examples stems from a more fundamental property: that they represent problems the base model cannot reliably solve.

We categorize training examples into four types based on base model performance (here base means the model pre GRPO training):

1. **Base Wrong:** Examples the base model gets wrong (accuracy <25% )
2. **Base Right:** Examples the base model gets right (accuracy  $\geq 25\%$  )
3. **Base Right, Same Size As Base Wrong:** A subset of the full Base Right Set that is the same size as Base Wrong, to account for model size
4. **All:** All the available examples in the train set

We hypothesize that training just on the examples where the base model is wrong will significantly match the base right subsets, as it provides the most opportunity for genuine capability improvement and higher learnable percentage.

#### 4.3.1 RESULTS

Table 5 presents the absolute performance after training on each subset:

Training Set	ShuffleObj			GSM8K		
	Phi-4	Qwen3-4B	Qwen3-8B	Phi-4	Qwen3-4B	Qwen3-8B
Base Wrong	<b>0.98</b>	<b>0.73</b>	0.81	<b>0.92</b>	<b>0.84</b>	<b>0.86</b>
Base Right	<b>0.98</b>	0.64	0.83	0.80	0.70	0.72
Base Right (Same Size)	<b>0.98</b>	0.67	0.75	0.78	0.68	0.71
All	0.97	0.66	<b>0.84</b>	0.85	0.77	0.85

Table 5: Absolute performance on test sets after training on different subsets based on base model performance. Bold indicates best or tied-best performance.

Table 6 shows the relative improvements of Base Wrong training over both Base Right subsets. Table 7 shows the sizes of each subset - Base Wrong is, on average, less than half the size of Base Right.

Metric	ShuffleObj			GSM8K			Average
	Phi-4	Qwen3-4B	Qwen3-8B	Phi-4	Qwen3-4B	Qwen3-8B	
<i>Improvement Over Base Right (Same Size)</i>							
Absolute	0.0%	6.0%	5.8%	14.0%	16.0%	15.0%	10.2%
Relative	0.0%	9.0%	7.7%	17.9%	23.5%	21.1%	14.3%
<i>Improvement Over Base Right (Full)</i>							
Absolute	0.0%	9.0%	-2.0%	12.0%	14.0%	14.0%	7.8%
Relative	0.0%	14.1%	-2.4%	15.0%	20.0%	19.4%	11.0%

Table 6: Improvements from training on Base Wrong examples compared to Base Right examples. Relative improvement calculated as (Base Wrong - Base Right) / Base Right.

#### 4.3.2 ANALYSIS

The results strongly support our hypothesis that GRPO learning value concentrates in examples the base model gets wrong:

**Base Wrong consistently outperforms Base Right:** Across models and tasks, training exclusively on Base Wrong examples achieves equal or superior performance to training on Base Right, even when Base Right has significantly more examples. The average relative improvement of 14.3% over size-matched Base Right subsets demonstrates that these examples provide fundamentally different learning signals. Even when we train on the full Base Right set, we see that Base Wrong outperforms it by 11% on average.

**Base Wrong matches / outperforms All:** Training on all Base Wrong samples actually slightly outperforms training on All samples on average (85% vs 82%), further supporting that most, if not all, meaningful training signal is concentrated in the Base Wrong subset.

**Task-specific patterns emerge:** On GSM8K, the advantage is particularly pronounced, with Base Wrong training achieving 15-23.5% relative improvements. This aligns with our earlier findings that mathematical reasoning benefits most from challenging examples. ShuffleObj shows more modest gains, potentially due to the discrete nature of object tracking limiting the diversity of solution paths.

**Model-specific sensitivity:** Phi-4 shows less differentiation between strategies on ShuffleObj (all achieving 0.98), suggesting potential ceiling effects. However, on GSM8K, even Phi-4 shows substantial 17.9% relative improvement, confirming the generality of the phenomenon.

**Practical implications:** These findings provide a practical recipe for GRPO dataset curation: practitioners can identify valuable training examples by simply evaluating base model performance, without needing to know which examples will ultimately improve. This avoids the circular dependency inherent in selecting “learnable” examples and provides a straightforward criterion for data collection.

Training Set	ShuffleObj			GSM8K		
	Phi-4	Qwen3-4B	Qwen3-8B	Phi-4	Qwen3-4B	Qwen3-8B
Base Wrong	17	37	13	343	409	371
Base Right	108	88	112	657	591	629
All	125	125	125	1000	1000	1000

Table 7: Subset sizes for each dataset, model, and training strategy. On average, Base Wrong is only 42% the size of Base Right, and yet consistently significantly outperforms this subset. Note that GSM8K is clipped to 1000 since we only run for 1000 steps with 1 prompt per update.

The concentration of learning value in Base Wrong examples validates our central thesis from a complementary angle. Our two experimental approaches converge on the same insight: examples with low pass@k rates (our “hardest” selection) are predominantly those where the base model fails (“base wrong”), while examples with high pass@k rates (our “easy” selection) are those the base model already masters (“base right”). This convergence demonstrates that whether we measure difficulty through multi-sample success rates or binary base model performance, we identify the same high-value training examples: those at the frontier of model capabilities where learning can actually occur.

## 5 CONCLUSION

This paper investigated a critical question for resource-constrained language model fine-tuning: under a fixed budget for example selection, which difficulty-based strategy maximizes the effectiveness of GRPO? Through comprehensive experiments across multiple models and tasks, we established that training on the hardest examples consistently yields superior performance on reasoning benchmarks.

We validated this “hard examples are all you need” principle from two complementary perspectives. First, our difficulty-based selection experiments showed that examples with the lowest pass@k success rates drive performance gains up to 47%, while easy examples yield minimal improvements. Second, our analysis of training data categorized by base model performance revealed that nearly all GRPO learning value concentrates in examples the base model initially fails to solve. Training exclusively on examples where the base model achieves low success rates ( $\hat{p}(x) < 0.25$ ) yields up

---

486 to 23.5% relative improvements compared to training on examples the base model already solves  
487 correctly ( $\hat{p}(x) \geq 0.25$ ).  
488

489 This convergence of evidence from both perspectives confirms a fundamental insight: GRPO learning  
490 requires training at the frontier of model capabilities. Examples the base model already masters  
491 reliably (high pass@k or those it solves correctly) are effectively the easiest and contribute minimal  
492 learning value, while examples where the model struggles (low pass@k or those it fails to solve)  
493 provide the variance in outcomes necessary for GRPO’s contrastive learning mechanism. This unified  
494 understanding provides practitioners with multiple practical approaches for identifying high-value  
495 training data while avoiding circular dependencies in dataset curation.

## 496 REFERENCES 497

498 Marah Abdin, Jyoti Aneja, Harkirat Behl, Sébastien Bubeck, Ronen Eldan, Suriya Gunasekar,  
499 Michael Harrison, Russell J. Hewett, Mojan Javaheripi, Piero Kauffmann, James R. Lee, Yin Tat  
500 Lee, Yuanzhi Li, Weishung Liu, Caio C. T. Mendes, Anh Nguyen, Eric Price, Gustavo de Rosa, Olli  
501 Saarikivi, Adil Salim, Shital Shah, Xin Wang, Rachel Ward, Yue Wu, Dingli Yu, Cyril Zhang, and  
502 Yi Zhang. Phi-4 technical report, 2024. URL <https://arxiv.org/abs/2412.08905>.

503 Mislav Balunović, Jasper Dekoninck, Ivo Petrov, Nikola Jovanović, and Martin Vechev. Math-  
504 arena: Evaluating llms on uncontaminated math competitions, February 2025. URL <https://matharena.ai/>.  
505

507 Paul F. Christiano, Jan Leike, Tom B. Brown, Miljan Martic, Shane Legg, and Dario Amodei. Deep  
508 reinforcement learning from human preferences. In Isabelle Guyon, Ulrike von Luxburg, Samy  
509 Bengio, Hanna M. Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman Garnett (eds.),  
510 *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information  
511 Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pp. 4299–4307, 2017.  
512

513 Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser,  
514 Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John  
515 Schulman. Training verifiers to solve math word problems. *arXiv:2110.14168*, 2021. URL  
516 <https://arxiv.org/abs/2110.14168>. Introduces GSM8K.

517 Ganqu Cui, Lifan Yuan, Zefan Wang, Hanbin Wang, Wendi Li, Bingxiang He, Yuchen Fan, Tianyu  
518 Yu, Qixin Xu, Weize Chen, Jiarui Yuan, Huayu Chen, Kaiyan Zhang, Xingtai Lv, Shuo Wang,  
519 Yuan Yao, Xu Han, Hao Peng, Yu Cheng, Zhiyuan Liu, Maosong Sun, Bowen Zhou, and Ning  
520 Ding. Process reinforcement through implicit rewards. *arXiv preprint arXiv:2502.01456*, 2025.  
521 doi: 10.48550/arXiv.2502.01456. URL <https://arxiv.org/abs/2502.01456>. PRIME:  
522 online process rewards from outcome labels.

523 Michael Han Daniel Han and Unsloth team. Unsloth, 2023. URL [http://github.com/  
524 unslothai/unsloth](http://github.com/unslothai/unsloth).  
525

526 Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Akhil Kadian, Ahmad Al-Dahle, et al. The  
527 llama 3 herd of models. *arXiv preprint arXiv:2407.21783*, 2024.  
528

529 Amirata Ghorbani and James Zou. Data shapley: Equitable valuation of data for machine learning. In  
530 *ICML, 2019*. URL <https://proceedings.mlr.press/v97/ghorbani19c.html>.  
531

532 Guy Hachohen and Daphna Weinshall. On the power of curriculum learning in training deep networks.  
533 In *ICML, 2019*. URL <https://arxiv.org/abs/1904.03626>.

534 Alex Havrilla, Edward Hughes, Mikayel Samvelyan, and Jacob Abernethy. Sparq: Synthetic problem  
535 generation for reasoning via quality-diversity algorithms, 2025. URL [https://arxiv.org/  
536 abs/2506.06499](https://arxiv.org/abs/2506.06499).  
537

538 Krishnateja Killamsetty, Durga Sivasubramanian, Ganesh Ramakrishnan, Abir De, and Rishabh Iyer.  
539 Grad-match: Gradient matching based data subset selection for efficient deep model training. In  
*ICML, 2021a*. URL <https://arxiv.org/abs/2103.00123>.

---

540 KrishnaTeja Killamsetty, Durga Sivasubramanian, Ganesh Ramakrishnan, and Rishabh K. Iyer.  
541 Glister: Generalization based data subset selection for efficient and robust learning. In *AAAI*,  
542 2021b. URL <https://ojs.aaai.org/index.php/AAAI/article/view/16988>.  
543

544 Zichen Liu, Changyu Chen, Wenjun Li, Penghui Qi, Tianyu Pang, Chao Du, Wee Sun Lee, and Min  
545 Lin. Understanding rl-zero-like training: A critical perspective. *arXiv preprint arXiv:2503.20783*,  
546 2025. doi: 10.48550/arXiv.2503.20783. URL <https://arxiv.org/abs/2503.20783>.  
547 Introduces Dr. GRPO; analyzes GRPO length bias.

548 OpenCompass. Aime2025 dataset. Hugging Face Hub, 2025. [https://huggingface.co/](https://huggingface.co/datasets/opencompass/AIME2025)  
549 [datasets/opencompass/AIME2025](https://huggingface.co/datasets/opencompass/AIME2025), accessed 2025-08-14.

550 Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong  
551 Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton,  
552 Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul F. Christiano, Jan Leike, and  
553 Ryan Lowe. Training language models to follow instructions with human feedback. In *NeurIPS*,  
554 2022. URL <https://arxiv.org/abs/2203.02155>.

555 Mansheej Paul, Surya Ganguli, and Gintare Karolina Dziugaite. Deep learning on a data diet: Finding  
556 important examples early in training. In *NeurIPS*, 2021. URL [https://arxiv.org/abs/](https://arxiv.org/abs/2107.07075)  
557 [2107.07075](https://arxiv.org/abs/2107.07075).

558 John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy  
559 optimization algorithms. *arXiv:1707.06347*, 2017. URL [https://arxiv.org/abs/1707.](https://arxiv.org/abs/1707.06347)  
560 [06347](https://arxiv.org/abs/1707.06347).

561 Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Mingchuan Zhang, Y. K. Li,  
562 Y. Wu, and Daya Guo. Deepseekmath: Pushing the limits of mathematical reasoning in open  
563 language models. *CoRR*, abs/2402.03300, 2024. doi: 10.48550/ARXIV.2402.03300. URL  
564 <https://doi.org/10.48550/arXiv.2402.03300>.  
565 <https://doi.org/10.48550/arXiv.2402.03300>.  
566

567 Yifan Sun, Jingyan Shen, Yibin Wang, Tianyu Chen, Zhendong Wang, Mingyuan Zhou, and Huan  
568 Zhang. Improving data efficiency for llm reinforcement fine-tuning through difficulty-targeted  
569 online data selection and rollout replay. *arXiv preprint arXiv:2506.05316*, 2025. doi: 10.48550/  
570 [arXiv.2506.05316](https://arxiv.org/abs/2506.05316). URL <https://arxiv.org/abs/2506.05316>.

571 Mirac Suzgun, Nathan Scales, Nathanael Schärli, Sebastian Gehrmann, Yi Tay, Hyung Won Chung,  
572 Aakanksha Chowdhery, Quoc V. Le, Ed H. Chi, Denny Zhou, and Jason Wei. Challenging big-bench  
573 tasks and whether chain-of-thought can solve them. In *Findings of ACL*, 2023. URL <https://aclanthology.org/2023.findings-acl.824.pdf>. BIG-Bench Hard; includes  
574 Tracking Shuffled Objects.  
575

576 Xiaoyu Tian, Sitong Zhao, Haotian Wang, Shuaiting Chen, Yiping Peng, Yunjie Ji, Han Zhao, and  
577 Xiangang Li. Deepdistill: Enhancing llm reasoning capabilities via large-scale difficulty-graded  
578 data training, 2025. URL <https://arxiv.org/abs/2504.17565>.  
579

580 Mariya Toneva, Alessandro Sordoni, Rémi Tachet des Combes, Adam Trischler, Yoshua Bengio,  
581 and Geoffrey J. Gordon. An empirical study of example forgetting during deep neural network  
582 learning. In *ICLR*, 2019. URL <https://arxiv.org/abs/1812.05159>.

583 Brandon Trabucco, Qasim Wani, Benjamin Pikus, and Vasu Sharma. Understanding trade offs when  
584 conditioning synthetic data, 2025. URL <https://arxiv.org/abs/2507.02217>.  
585

586 Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc V. Le, Ed H. Chi, Sharan Narang, Aakanksha  
587 Chowdhery, and Denny Zhou. Self-consistency improves chain of thought reasoning in language  
588 models. *arXiv:2203.11171*, 2022. URL <https://arxiv.org/abs/2203.11171>.

589 An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang  
590 Gao, Chengen Huang, Chenxu Lv, Chujie Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng Hu,  
591 Hao Ge, Haoran Wei, Huan Lin, Jialong Tang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin  
592 Yang, Jiayi Yang, Jing Zhou, Jingren Zhou, Junyang Lin, Kai Dang, Keqin Bao, Kexin Yang,  
593 Le Yu, Lianghao Deng, Mei Li, Mingfeng Xue, Mingze Li, Pei Zhang, Peng Wang, Qin Zhu, Rui  
Men, Ruize Gao, Shixuan Liu, Shuang Luo, Tianhao Li, Tianyi Tang, Wenbiao Yin, Xingzhang

---

594 Ren, Xinyu Wang, Xinyu Zhang, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yinger  
595 Zhang, Yu Wan, Yuqiong Liu, Zekun Wang, Zeyu Cui, Zhenru Zhang, Zhipeng Zhou, and Zihan  
596 Qiu. Qwen3 technical report, 2025. URL <https://arxiv.org/abs/2505.09388>.  
597

## 598 A RELATED WORK 599

600 Our work intersects several research areas: reinforcement learning from human feedback, example  
601 selection and curriculum learning, reasoning benchmarks, and knowledge evaluation tasks. We review  
602 each area and position our contributions within this broader context.  
603

### 604 A.1 REINFORCEMENT LEARNING FROM HUMAN FEEDBACK AND GRPO 605

606 Reinforcement learning from human feedback (RLHF) has emerged as a dominant paradigm for  
607 aligning language models with human preferences Christiano et al. (2017); Ouyang et al. (2022). The  
608 standard RLHF pipeline involves training a reward model from human preferences, then optimizing  
609 the language model policy using algorithms like Proximal Policy Optimization (PPO) Schulman et al.  
610 (2017). However, PPO-based RLHF suffers from high computational costs, training instability, and  
611 significant variance in gradient estimates.

612 Group Relative Policy Optimization (GRPO), introduced in DeepSeekMath Shao et al. (2024),  
613 addresses these challenges by reformulating the optimization problem. Instead of training a separate  
614 reward model, GRPO directly optimizes using group-relative advantages computed from multiple  
615 samples per prompt. For each training prompt, GRPO samples a group of  $G$  outputs, computes their  
616 rewards (e.g., correctness on math problems), and uses the deviation from the group mean as the  
617 advantage signal. This approach reduces variance through the group baseline while eliminating the  
618 need for reward model training.

619 Our work builds directly on GRPO’s formulation but introduces a critical new dimension: budget-  
620 aware example selection. While previous GRPO applications assume uniform sampling from the  
621 training distribution, we show that strategic selection based on example difficulty can substantially  
622 improve learning efficiency, particularly relevant when data acquisition costs constrain the training  
623 set size.  
624

### 625 A.2 EXAMPLE DIFFICULTY AND DATA SELECTION 626

627 The importance of example selection in machine learning has long been recognized, with multiple  
628 research threads addressing different aspects of this problem.

629 **Curriculum learning** suggests that training on examples in order of increasing difficulty can improve  
630 convergence and final performance Hachohen & Weinshall (2019). However, curriculum learning  
631 typically assumes access to the full dataset and focuses on presentation order rather than subset  
632 selection. Our work differs by operating under a strict budget constraint where only a fraction of  
633 examples can be used.

634 **Data valuation and influence** methods aim to identify the most valuable training examples. Data  
635 Shapley Ghorbani & Zou (2019) uses cooperative game theory to assign value to each training  
636 example based on its marginal contribution. However, these methods are computationally expensive  
637 and designed for supervised learning rather than RLHF settings.

638 **Coreset selection** techniques identify representative subsets that approximate the full dataset’s  
639 gradient Killamsetty et al. (2021a;b). GradMatch selects examples whose gradients best match the  
640 full dataset gradient, while GLISTER focuses on generalization-based selection. These methods  
641 provide theoretical guarantees but require access to labels and gradient computation that may not be  
642 available in our budget-constrained setting.

643 **Example forgetting and memorization** studies have shown that neural networks exhibit consistent  
644 patterns in which examples are learned or forgotten during training Toneva et al. (2019).  
645

646 **Synthetic data** is often an area of investigation for training, especially since it is often easier to  
647 control for difficulty. Our work doesn’t explicitly investigate using synthetic data and instead selects  
from a real, static larger dataset Trabucco et al. (2025); Tian et al. (2025); Havrilla et al. (2025).

“Unforgettable” examples are learned early and retained, while others are repeatedly learned and forgotten. The Data Diet work Paul et al. (2021) leverages Error L2-Norm (EL2N) scores computed early in training to identify important examples, showing that networks can be trained on small subsets without performance loss.

Our approach differs from these prior works in several ways: (1) We operate in the RLHF/GRPO setting where rewards are binary and computed at inference time; (2) We use the base model’s multi-sample success rate as a difficulty proxy, requiring no training to compute; (3) We explicitly compare different difficulty-based selection strategies under identical budgets.

### A.3 REASONING BENCHMARKS AND EVALUATION

Evaluating reasoning capabilities in language models requires carefully designed benchmarks that test different aspects of logical and mathematical thinking.

**GSM8K** Cobbe et al. (2021) has become the standard benchmark for grade-school mathematical reasoning, containing 8,792 problems requiring multi-step arithmetic and logical reasoning. The dataset’s problems are linguistically diverse but mathematically elementary, making it ideal for studying reasoning without confounding advanced mathematical knowledge.

**BIG-Bench Hard (BBH)** Suzgun et al. (2023) curates 23 challenging tasks from the broader BIG-Bench suite where language models initially showed poor performance. The Tracking Shuffled Objects task, which we use in our experiments, requires models to track entity positions through a series of swapping operations: a pure reasoning task with minimal knowledge requirements.

**Chain-of-thought prompting and self-consistency** Wang et al. (2022) have proven effective for improving reasoning performance. Self-consistency samples multiple reasoning paths and aggregates answers, often yielding substantial accuracy improvements. We leverage this insight in our difficulty estimation, using multi-sample success rates to robustly measure example hardness.

### A.4 GRPO

The GRPO algorithm operates by sampling groups of outputs for each training prompt and using relative performance within each group to compute advantages. Specifically, for each prompt  $q$  in the selected subset, the algorithm samples  $G$  outputs  $\{o_1, \dots, o_G\}$  from the current policy  $\pi_\theta$ , computes binary rewards  $r_i \in \{0, 1\}$  based on correctness, and calculates a group baseline  $\bar{r} = \frac{1}{G} \sum_{i=1}^G r_i$ . The advantage for each output is then computed as  $A_i = r_i - \bar{r}$ , which naturally centers the learning signal around the group’s average performance. These advantages drive policy updates while KL regularization to a reference policy  $\pi_{\text{ref}}$  prevents catastrophic distribution shifts.

The training objective maximizes the expected advantage-weighted log-likelihood while minimizing KL divergence from the reference policy:

$$\mathcal{L}_{\text{GRPO}} = -\mathbb{E}_{q \sim S} \mathbb{E}_{o \sim \pi_\theta(\cdot|q)} [A(o, q) \cdot \log \pi_\theta(o|q) - \beta \cdot \text{KL}(\pi_\theta || \pi_{\text{ref}})]$$

where  $\beta$  controls the strength of KL regularization.

### A.5 CLOSEST GRPO WORK

Recent work has explored various aspects of GRPO optimization that relate closely to our budget-aware selection study:

**Difficulty-targeted selection and replay.** Sun et al. propose a data-efficiency recipe for GRPO-style RL fine-tuning that (i) targets adaptive difficulty online, prioritizing questions of moderate difficulty expected to yield informative gradients, and (ii) introduces rollout replay to reuse recent samples and lower per-step cost. Across six LLM-dataset pairs they report 25-65% faster time-to-target performance than vanilla GRPO, suggesting that what you train on (and when) can matter as much as how much you train. Our paper shares the focus on difficulty-aware allocation under constraints, but differs in scope: we study budgeted, offline subset selection (hard vs. medium vs. easy vs. random) with a fixed labeling budget and analyze why “hardest” can dominate on reasoning tasks, whereas Sun et al. optimize online time-to-quality with an adaptive, typically moderate-difficulty target and replay Sun et al. (2025).

**Process rewards compatible with GRPO.** Cui et al. (PRIME) tackle the sparsity and credit-assignment limits of outcome-only rewards by learning implicit process reward models online from outcome labels alone, then combining dense token-level signals with outcome rewards during RL. PRIME is compatible with standard advantage formulations (including GRPO) and removes a separate reward-modeling stage; empirically, they show sizable gains on competitive math and coding (e.g., 15% average over SFT with a 7B base). Our study focuses on which examples to buy under an outcome-reward GRPO setup; PRIME is complementary, and our framework could be extended to ask whether difficulty-aware selection interacts with dense process rewards in similar ways (e.g., whether “hardest” remains optimal when per-step feedback reduces variance) Cui et al. (2025).

**GRPO biases and token efficiency.** Liu et al. present a critical analysis of R1-Zero-like training and identify an optimization bias in GRPO: the objective can inflate response length, particularly for incorrect outputs, harming token efficiency. They introduce Dr. GRPO, an unbiased variant that preserves reasoning quality while improving efficiency, and demonstrate strong AIME-24 results with a 7B base model. For our budget-aware study, their findings motivate length-controlled evaluations and guardrails when comparing difficulty-conditioned subsets (e.g., ensuring that harder-subset gains are not confounded by pathological length growth and considering Dr. GRPO as a robustness check) Liu et al. (2025).

**Positioning.** Together, these works closest to GRPO underscore three levers we explicitly separate and test: (a) example selection by difficulty (Sun et al.), (b) reward shaping via dense process feedback (Cui et al.), and (c) objective design to avoid length bias (Liu et al.). Our contribution is orthogonal and complementary: a budget-aware, offline selection protocol and theory-informed analysis showing why, under outcome-level GRPO, training on the hardest decile can sustain a higher fraction of “learnable” examples (non-degenerate group variance) and thus larger gains on reasoning tasks, while we adopt length controls in line with Dr. GRPO insights Sun et al. (2025); Cui et al. (2025); Liu et al. (2025).

## A.6 POSITIONING OUR CONTRIBUTIONS

Our work makes several distinct contributions relative to this prior literature:

1. We are the first to systematically study budget-aware example selection specifically for GRPO fine-tuning, addressing a practical constraint faced by practitioners.
2. We provide a unified evaluation framework comparing difficulty-based selection strategies across multiple task types and model scales.
3. We offer theoretical insight connecting GRPO’s group-advantage mechanism to the effectiveness of hard-example selection, grounded in variance reduction principles from reinforcement learning theory.
4. We demonstrate that selection strategy effectiveness depends critically on task type, with reasoning tasks benefiting from hard examples while knowledge tasks show no such preference.

These contributions provide both theoretical understanding and practical guidance for efficient RLHF fine-tuning under resource constraints.

## B HYPERPARAMETERS

We set the group size  $G = 8$  for our main experiments based on preliminary studies showing this balances variance reduction with computational efficiency. The KL coefficient  $\beta = 0.1$  provides sufficient regularization without overly constraining learning. Training proceeds for 1000 steps with a learning rate of  $3 \times 10^{-5}$  using cosine decay, processing 1 prompt per gradient update due to memory constraints. We use 4-bit quantized models from Unsloth Daniel Han & team (2023). We evaluate model performance every 100 steps to track learning dynamics, using the initial SFT checkpoint as our reference policy throughout training.

756  
757  
758  
759  
760  
761  
762  
763  
764  
765  
766  
767  
768  
769  
770  
771  
772  
773  
774  
775  
776  
777  
778  
779  
780  
781  
782  
783  
784  
785  
786  
787  
788  
789  
790  
791  
792  
793  
794  
795  
796  
797  
798  
799  
800  
801  
802  
803  
804  
805  
806  
807  
808  
809

---

## C LLM USE

Large Language Models (LLMs) were used for the following purposes during this research:

1. **Writing some of the code for the experiments:** LLMs assisted in generating and debugging portions of the experimental code, including data processing scripts and evaluation utilities.
2. **Proofreading the text in the paper:** LLMs were used to identify grammatical errors, improve clarity, and ensure consistency in the manuscript.
3. **LaTeX support:** LLMs provided assistance with LaTeX syntax for complex table formatting and figure placement.