

Placing Puzzle Pieces Where They Matter: A Question Augmentation Framework for Reinforcement Learning

Yangyi Fang¹, Haolin Shi^{1,†}

¹Tsinghua University shihaolin0720@gmail.com

Abstract

Reinforcement learning has become a powerful approach for enhancing large language model reasoning, but faces a fundamental dilemma: training on easy problems can cause overfitting and pass@k degradation, while training on hard problems often results in sparse rewards. Recent question augmentation methods address this by prepending partial solutions as hints. However, uniform hint provision may introduce redundant information while missing critical reasoning bottlenecks, and excessive hints can reduce reasoning diversity, causing pass@k degradation. We propose **PieceHint**, a hint injection framework that strategically identifies and provides critical reasoning steps during training. By scoring the importance of different reasoning steps, selectively allocating hints based on problem difficulty, and progressively withdrawing scaffolding, PieceHint enables models to transition from guided learning to independent reasoning. Experiments on six mathematical reasoning benchmarks show that our 1.5B model achieves comparable average performance to 32B baselines while preserving pass@k diversity across all k values.

1 Introduction

Reinforcement learning (RL) has emerged as a powerful paradigm for enhancing LLM reasoning capabilities (Guo et al., 2025a; Moshkov et al., 2025). However, RL-based training faces a fundamental dilemma: training on easy problems causes overfitting and pass@k degradation, while training on hard problems yields sparse reward signals (Yue et al., 2025a; Shao et al., 2024).

Various approaches address this challenge through reward shaping (Zhu et al., 2025; Shao et al., 2025), process-level supervision (Wang et al., 2024; Malik et al., 2025), and curriculum learning (Luo et al., 2025; Parashar et al., 2025;

[†] Corresponding author.

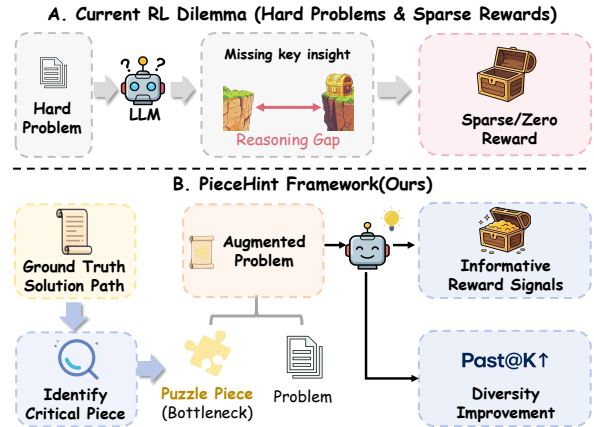


Figure 1: RL training dilemma and PieceHint solution. **Top:** Missing key insights at reasoning bottlenecks lead to sparse/zero rewards. **Bottom:** PieceHint identifies critical pieces from ground-truth solutions, injects them as hints to bridge gaps, enabling informative rewards and pass@k diversity.

Yan et al., 2025b; Li et al., 2025). Question augmentation offers a complementary direction: StepHint (Zhang et al., 2025) prepends partial ground-truth solutions as hints, enabling models to discover successful trajectories on hard problems. However, uniform hint provision may introduce redundant information while missing critical bottlenecks, and excessive hints can reduce reasoning diversity, causing pass@k degradation.

Figure 1 illustrates the core challenge and our solution. In reinforcement learning, models receive rewards only upon generating correct final answers—partial progress yields no credit. When facing hard problems, models often fail at specific reasoning bottlenecks—missing a pivotal lemma, overlooking a non-obvious insight, or lacking a critical technique. These bottlenecks act as missing puzzle pieces: without them, the reasoning chain cannot be completed. If all attempts fail at these bottlenecks, reward variance becomes zero, eliminating gradient signals for learning (Figure 1, top).

Uniform hint provision may address routine derivations the model can already handle, while leaving critical reasoning gaps unbridged. This reveals a key insight: reasoning breakthroughs occur at specific bottlenecks rather than uniformly across steps, and identifying and providing these critical puzzle pieces yields maximum learning benefit.

To address this challenge, we propose **PieceHint**, which treats critical reasoning steps as puzzle pieces that unlock solution pathways (Figure 1, bottom). Our approach strategically identifies high-value pieces through value-driven scoring, allocates them based on problem difficulty and model capability, and progressively withdraws them during training, enabling models to transition from scaffolded learning to independent reasoning.

Our contributions are:

- We propose PieceHint, a hint injection framework that identifies critical reasoning pieces and strategically provides them during training.
- We introduce progressive piece withdrawal curriculum, enabling transition from scaffolded to hint-free reasoning.
- Experiments on six mathematical benchmarks show our 1.5B model achieves comparable average performance to 32B baselines while demonstrating superior parameter efficiency and preserving exploration diversity.

2 Related Works

Reinforcement Learning for Reasoning. Reinforcement learning has emerged as a powerful paradigm for enhancing LLM reasoning, particularly in mathematical problem-solving (Guo et al., 2025a; Moshkov et al., 2025). PPO (Schulman et al., 2017) and GRPO (Guo et al., 2025a) establish the algorithmic basis for policy optimization with outcome supervision. Subsequent work has explored improvements across multiple dimensions: reward shaping (Zhu et al., 2025; Shao et al., 2025), process-level supervision (Wang et al., 2024; Malik et al., 2025), and chain-of-thought verification (Wen et al., 2025; Liu et al., 2023; Liu, 2024; Liu et al., 2025b; Li et al., 2026).

Sample Efficiency and Exploration. Parallel efforts address sample efficiency and exploration through uncertainty-aware training strategies.

TreeRL (Hou et al., 2025) and VinePPO (Kazemnejad et al., 2025) leverage entropy or confidence metrics to prioritize high-quality rollouts and improve credit assignment. MRT (Qu et al., 2025) reuses partial trajectories at inference time to amortize computational cost across multiple queries. Curriculum learning strategies have also been proposed to organize training progression, either by solution length (Luo et al., 2025; Yu et al., 2025b), problem difficulty (Parashar et al., 2025), or reasoning complexity (Yan et al., 2025b).

The Dilemma of Problem Difficulty in RL Training. Recent work (Shao et al., 2024; Yue et al., 2025a; Zhao et al., 2025) reveals a fundamental tension in RL-based reasoning: training on easy problems can lead to pass@k degradation as models may overfit to familiar solution patterns, while training on hard problems often results in sparse reward signals. To mitigate pass@k degradation, several methods (Yu et al., 2025a; Liu et al., 2025a; Fang et al., 2025) employ entropy regularization through adaptive clipping, temperature tuning, or periodic reference model updates. These algorithmic modifications operate at the optimization level, adjusting how models learn from existing rollouts.

Question Augmentation for RL. Question augmentation operates at the input level to reshape problem difficulty. StepHint (Zhang et al., 2025) prepends sequential partial solutions (e.g., the first 50% of reasoning steps), bridging the gap between easy and hard problems while improving convergence speed and pass@k preservation. An alternative approach (Yan et al., 2025a) provides entire reasoning chains from stronger models as reference trajectories, though the target model still explores independently most of the time. While these methods demonstrate effectiveness, position-based selection may include routine steps while potentially missing non-obvious insights that appear later in solutions, as validated by our ablation studies (Section 5.3). In contrast to position-based approaches, our method identifies critical reasoning bottlenecks through value-driven piece identification and scoring, strategically providing hints based on their impact on learning rather than sequential position. Furthermore, through progressive piece withdrawal, we enable models to transition from scaffolded learning to hint-free reasoning, addressing potential generalization limitations of fixed augmentation methods.

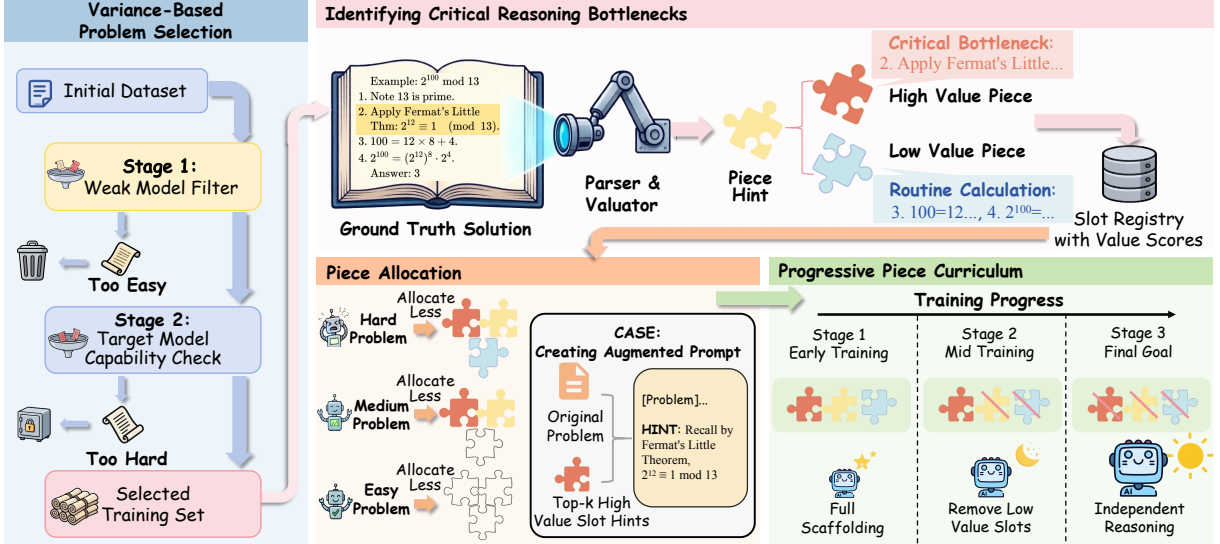


Figure 2: Overview of PieceHint. We identify critical reasoning pieces via value scores, allocate them by model capability, and progressively withdraw hints during training for hint-free reasoning.

3 Preliminaries

Problem Formulation. We consider mathematical reasoning as a sequence generation task. Given a problem q from problem space \mathcal{Q} , the goal is to generate a reasoning trajectory o that leads to the correct answer. We denote the ground-truth solution as $\tau^* = [\tau_1^*, \tau_2^*, \dots, \tau_n^*]$, which is a sequence of reasoning steps, and the final answer as a^* . A generated response o is considered correct if its extracted answer matches a^* .

In question augmentation, we construct augmented prompts by prepending auxiliary information as hints to the original problem. Formally, an augmented prompt is denoted as $\tilde{q} = [h; q]$, where h represents the hint and $[\cdot; \cdot]$ denotes concatenation. The hint h can take various forms depending on the augmentation strategy—for instance, partial solution steps, key intermediate results, or critical reasoning components. During RL training, the model learns from augmented prompts \tilde{q} while being evaluated on original prompts q without any hints, ensuring the learned policy generalizes to hint-free scenarios.

Group Relative Policy Optimization. We adopt Group Relative Policy Optimization (GRPO) (Shao et al., 2024) as our reinforcement learning algorithm. GRPO addresses the high variance problem in policy gradient methods by sampling multiple rollouts per prompt and normalizing advantages within each group.

Formally, for each prompt q , GRPO samples a group of G rollouts $\{o_i\}_{i=1}^G$ from the current policy

π_{old} . The advantage function for the i -th rollout is computed via group-wise normalization:

$$A_{i,t} = \frac{R(o_i) - \bar{R}}{\sigma_R}, \quad (1)$$

where $R(o_i)$ denotes the reward for response o_i , and \bar{R} and σ_R represent the mean and standard deviation of rewards within the group $\{R(o_j)\}_{j=1}^G$, respectively. This advantage $A_{i,t}$ is shared across all tokens in the sequence.

The optimization objective maximizes a clipped surrogate loss:

$$\mathcal{L}(\theta) = \mathbb{E}_{q \sim \mathcal{D}, \{o_i\} \sim \pi_{\text{old}}} \left[\frac{1}{\sum_i |o_i|} \sum_{i=1}^G \sum_{t=1}^{|o_i|} \min \left(r_{i,t} A_{i,t}, \text{clip}(r_{i,t}, 1 - \epsilon, 1 + \epsilon) A_{i,t} \right) \right], \quad (2)$$

where $r_{i,t} = \pi_{\theta}(o_{i,t} | q, o_{i,<t}) / \pi_{\text{old}}(o_{i,t} | q, o_{i,<t})$ is the importance sampling ratio, ϵ is the clipping parameter, and $|o_i|$ denotes the length of the i -th response. Following recent work (Yu et al., 2025a; Liu et al., 2025c), we omit the KL divergence penalty term in our implementation.

We employ a binary reward function $R : \mathcal{Q} \times \mathcal{V}^* \rightarrow \{0, 1\}$, where $R(q, o_i) = 1$ if the final answer is correct and properly formatted, and $R(q, o_i) = 0$ otherwise. This sparse reward design provides clear credit assignment but can lead to insufficient learning signals on extremely difficult problems where success rates approach zero.

4 Methodology

The key insight driving our approach is that reasoning steps contribute unequally to problem-solving success: some represent critical bottlenecks—pivotal lemmas, non-obvious insights, or key techniques—where models frequently fail, while others are routine derivations models can handle independently. This motivates PieceHint, a value-driven framework that strategically provides hints at high-impact bottlenecks rather than uniformly covering all steps.

Our method operates through a pipeline combining offline preprocessing and online training (Figure 2). In the offline phase, we apply variance-based problem selection (Section 4.1) to identify problems in the optimal difficulty range from initial dataset $\mathcal{Q}_{\text{init}}$, yielding training set $\mathcal{Q}_{\text{train}}$. For each problem $q \in \mathcal{Q}_{\text{train}}$, value-driven piece identification (Section 4.2) extracts critical reasoning pieces from ground-truth solutions and assigns importance scores, followed by piece allocation (Section 4.3) that determines initial hint budgets based on model capability. During online training, progressive piece curriculum (Section 4.4) gradually withdraws hints, enabling transition from scaffolded learning to hint-free reasoning.

4.1 Variance-Based Problem Selection

Not all problems are equally suitable for reinforcement learning. Trivially easy problems cause overfitting and degrade exploration, while prohibitively difficult problems provide insufficient reward signals (Shao et al., 2024; Yue et al., 2025a). We propose a two-stage selection strategy identifying problems in the optimal difficulty range—challenging yet learnable.

Identifying hard problems. Given initial dataset $\mathcal{Q}_{\text{init}}$, we employ a weak selection model μ_{weak} (comparable or slightly weaker than our training model) to filter trivially easy problems. For each problem $q \in \mathcal{Q}_{\text{init}}$, we sample m reasoning attempts from μ_{weak} , counting correct solutions as $c_{\text{weak}}(q)$, and construct the hard problem subset:

$$\mathcal{Q}_{\text{hard}} = \{q \in \mathcal{Q}_{\text{init}} \mid c_{\text{weak}}(q) \leq \alpha_1 \cdot m\}, \quad (3)$$

where α_1 is a threshold parameter, retaining problems where μ_{weak} succeeds at most $\alpha_1 \cdot m$ times out of m attempts.

Capability-aligned selection. Among $\mathcal{Q}_{\text{hard}}$, we select problems aligning with our training model μ_0 ’s current capability. We evaluate μ_0 on each

problem without augmentation, sampling m attempts and computing $\text{pass}@m(q, \mu_0) = c_{\text{train}}(q)$ (the number of correct solutions out of m attempts), yielding:

$$\mathcal{Q}_{\text{train}} = \{q \in \mathcal{Q}_{\text{hard}} \mid \alpha_2 \cdot m \leq c_{\text{train}}(q) \leq \alpha_3 \cdot m\}, \quad (4)$$

where α_2 and α_3 bound the acceptable success rate range.

This two-stage filtering removes problems that are too easy ($c_{\text{weak}}(q) > \alpha_1 \cdot m$), intractable ($c_{\text{train}}(q) = 0$), or already accessible ($c_{\text{train}}(q) > \alpha_3 \cdot m$), retaining those with moderate success rates—the ideal regime for effective RL. We preserve the success counts $\{c_{\text{train}}(q)\}_{q \in \mathcal{Q}_{\text{train}}}$ for subsequent piece allocation in Section 4.3.

4.2 Identifying Critical Reasoning Bottlenecks

Having identified learnable problems, we next determine which solution parts are most critical. Traditional methods mechanically truncate solution prefixes; we propose a value-driven approach identifying critical bottlenecks where targeted hints provide maximum benefit.

For each problem $q \in \mathcal{Q}_{\text{train}}$ with ground-truth solution $\tau^* = [\tau_1^*, \dots, \tau_n^*]$ (defined in Section 3), we employ a stronger LLM to parse τ^* into a set of reasoning pieces $\mathcal{P}(q) = \{p_1, \dots, p_n\}$. Each piece p_i encapsulates a logically complete reasoning component that contributes to the solution—for instance, recognizing a key mathematical structure, establishing an intermediate lemma, applying a pivotal calculation technique, or making a critical variable substitution. The granularity of parsing is designed to capture semantic reasoning units rather than arbitrary sentence boundaries, ensuring each piece represents a meaningful cognitive step in the problem-solving process.

Value scoring for bottleneck identification. We employ a stronger LLM to assign each piece a discrete value score $V(p_i)$ based on three criteria: (1) *novelty*—whether the insight is non-trivial, (2) *difficulty*—whether deriving this step demands mathematical maturity, and (3) *impact*—whether subsequent reasoning critically depends on this result. High-value pieces correspond to critical bottlenecks where hints substantially reduce problem difficulty (scoring prompt details in Appendix I).

To enable cross-problem comparison, we normalize scores within each problem via min-max normalization:

$$V_{\min} = \min_{p_j \in \mathcal{P}(q)} V(p_j), \quad V_{\max} = \max_{p_j \in \mathcal{P}(q)} V(p_j), \quad (5)$$

$$D(p_i) = \frac{V(p_i) - V_{\min}}{V_{\max} - V_{\min}}, \quad (6)$$

yielding normalized values $D(p_i) \in [0, 1]$ with higher values indicating more critical bottlenecks. In practice, we observe that value distributions within problems exhibit clear stratification, with a small subset of high-value pieces (typically 1-3 per problem) standing out as pivotal reasoning steps. We maintain a global registry $\mathcal{R} = \{(q, \mathcal{P}(q), V, D)\}_{q \in \mathcal{Q}_{\text{train}}}$ storing piece structures and values for all training problems.

4.3 Piece Allocation

Having identified and scored pieces, we determine how many hints each problem receives. Problems vary in difficulty—some require substantial scaffolding, others benefit from minimal hints preserving exploration. We allocate initial piece budgets based on model capability.

Leveraging the success counts $\{c_{\text{train}}(q)\}_{q \in \mathcal{Q}_{\text{train}}}$ from Section 4.1, we assign initial piece budgets via a three-tier allocation function. For conciseness, we denote $c(q) = c_{\text{train}}(q)$ as the number of correct solutions out of m attempts. The allocation function is defined as:

$$k^{(0)}(q) = \begin{cases} k_{\max} & \text{if } c(q) \in [0, \beta_1 \cdot m], \\ \lfloor k_{\max}/2 \rfloor & \text{if } c(q) \in (\beta_1 \cdot m, \beta_2 \cdot m], \\ 0 & \text{if } c(q) \in (\beta_2 \cdot m, m], \end{cases} \quad (7)$$

where $k^{(0)}(q)$ denotes the initial number of pieces allocated to problem q , and $k_{\max}, \beta_1, \beta_2$ are hyperparameters controlling the allocation. Specifically, hard problems ($c(q) \in [0, \beta_1 \cdot m]$) receive maximum support (k_{\max} pieces) enabling successful trajectories; medium problems ($c(q) \in (\beta_1 \cdot m, \beta_2 \cdot m]$) receive moderate support ($\lfloor k_{\max}/2 \rfloor$ pieces) improving consistency; easy problems ($c(q) \in (\beta_2 \cdot m, m]$) receive no hints preserving exploration.

Value-driven piece selection. For problems with $k^{(0)}(q) > 0$, we select the top- $k^{(0)}(q)$ pieces maximizing the sum of normalized values:

$$H_{k^{(0)}}(q) = \arg \max_{\substack{S \subseteq \mathcal{P}(q) \\ |S|=k^{(0)}(q)}} \sum_{p_i \in S} D(p_i). \quad (8)$$

This yields an augmented training set $\{(q, H_{k^{(0)}}(q))\}_{q \in \mathcal{Q}_{\text{train}}}$, where $H_{k^{(0)}}(q)$ denotes the initial hint set for problem q . During training, pieces in $H_{k^{(0)}}(q)$ are prepended to the original problem to construct augmented prompts $\tilde{q} = [H_{k^{(0)}}(q); q]$.

4.4 Progressive Piece Curriculum

Having constructed augmented prompts with hints, we ensure the model eventually learns to solve problems independently. Fixed augmentation perpetually providing hints may hinder generalization; we implement a progressive curriculum gradually withdrawing hints during training, enabling transition from scaffolded learning to independent reasoning.

Using the augmented training set $\{(q, H_{k^{(0)}}(q))\}_{q \in \mathcal{Q}_{\text{train}}}$ from Section 4.3, we train via GRPO while implementing this curriculum. At each training iteration, we sample problems from $\mathcal{Q}_{\text{train}}$ and construct augmented prompts by prepending the current hint set $H_k(q)$ (initially set to $H_{k^{(0)}}(q)$). The model generates responses and receives binary rewards per Section 3, with policy updates following the GRPO objective. As training progresses, different problems advance through the curriculum at different rates depending on their sampling frequency. Each problem q maintains a sampling counter $s(q)$ (initialized to zero) tracking its training exposure.

Frequency-based withdrawal mechanism. When $s(q) \bmod N_{\text{check}} = 0$ (i.e., every N_{check} samples of problem q) and the current hint set is non-empty ($|H_k(q)| > 0$), we remove the least valuable piece to preserve critical bottlenecks longer:

$$p_{\text{remove}} = \arg \min_{p_i \in H_k(q)} D(p_i), \quad (9)$$

$$H_k(q) \leftarrow H_k(q) \setminus \{p_{\text{remove}}\},$$

where $D(p_i)$ is the normalized value from Eq. 6. This implements three design principles: (1) per-problem sampling frequency ensures proportional progression regardless of global training dynamics, (2) removing one piece per N_{check} samples enables smooth transitions—a problem with $k^{(0)}(q)$ pieces receives $k^{(0)}(q) \times N_{\text{check}}$ training samples with gradually decreasing support, and (3) removing low-value pieces first preserves support at critical bottlenecks while building independence on simpler steps.

For a problem initially assigned $k^{(0)}(q)$ pieces, all hints are withdrawn after exactly $k^{(0)}(q) \times N_{\text{check}}$ samples of that specific problem. Once $H_k(q) = \emptyset$, the problem continues training without any hints, forcing the model to solve it independently. This ensures the model develops robust hint-free reasoning capabilities that generalize to evaluation scenarios where no ground-truth information is available.

Table 1: Main results on mathematical reasoning benchmarks. We report pass@1 accuracy (%). The best results are **bolded**, and the second-best results are underlined.

Method	AIME24	AIME25	AMC23	MATH500	Minerva	Olympiad	Avg.
<i>Baseline Models</i>							
<i>1.5B Models</i>							
DeepSeek-R1-Distill-1.5B (Guo et al., 2025a)	24.4	21.8	63.6	72.1	26.2	43.4	41.9
Qwen3-1.7B (Yang et al., 2025)	30.6	47.5	69.0	84.1	35.4	53.7	53.4
Nemotron-1.5B (Moshkov et al., 2025)	44.0	34.2	82.7	88.5	22.3	61.2	55.5
<i>4-8B Models</i>							
DeepSeek-R1-Distill-7B (Guo et al., 2025a)	39.6	29.9	78.8	80.3	40.0	55.1	54.0
Qwen3-4B (Yang et al., 2025)	47.5	37.3	80.0	85.1	<u>40.4</u>	58.5	58.1
<i>32B Models</i>							
DeepSeek-R1-Distill-32B (Guo et al., 2025a)	<u>53.1</u>	<u>39.5</u>	<u>87.3</u>	<u>86.0</u>	45.1	<u>61.2</u>	<u>62.0</u>
<i>PieceHint (Ours)</i>							
PieceHint-DeepSeek-1.5B	35.8	25.5	71.2	87.7	36.7	50.1	51.2
PieceHint-Qwen3-1.7B	42.2	38.8	80.4	89.8	41.6	62.4	59.2
PieceHint-Nemotron-1.5B	54.5	43.7	89.1	91.3	28.5	68.3	62.6

5 Experiments

5.1 Experimental Setup

Datasets. We construct training datasets through a systematic filtering pipeline applied to the OpenR1-Math-220K dataset. Following the variance-based problem selection strategy described in Section 4.1, we employ DeepSeek-R1-Distill-1.5B as the weak selection model μ_{weak} to filter trivially easy problems. Each base model then uses its own initial checkpoint as the capability reference model μ_0 to perform capability-aligned selection, resulting in model-specific training sets $\mathcal{Q}_{\text{train}}$ containing problems within the target difficulty range for reinforcement learning.

Implementation Details. Our implementation is based on the VeRL training framework (Sheng et al., 2025). We configure the training pipeline with a rollout batch size of 512 and an update batch size of 32, performing 16 rollout samples per prompt. During rollout generation, we employ nucleus sampling with Top-p $p = 1.0$ and temperature $\tau = 1.0$. We train for 2,400 gradient update steps, corresponding to approximately 1 million rollout samples. All experiments are conducted on $8 \times$ NVIDIA A100 80GB GPUs. Full hyperparameters are provided in Appendix B.

Evaluation. We evaluate our method on six mathematical reasoning benchmarks: AIME24 (MAA, 2025), AIME25 (MAA, 2025), AMC23 (MAA, 2023), MATH500 (Hendrycks et al., 2021), Min-

erva (Lewkowycz et al., 2022), and Olympiad (He et al., 2024). During evaluation, we use Top-p sampling with $p = 1.0$ and temperature $\tau = 1.0$. We report pass@ k metrics following (Yue et al., 2025b) and (Chen et al., 2021), with answer verification performed using the Math-Verify tool (HuggingFace, 2025).

5.2 Main Results

Table 1 compares PieceHint against baseline models across six mathematical reasoning benchmarks. PieceHint-Nemotron-1.5B achieves competitive performance with the 32B DeepSeek baseline, reaching comparable average accuracy while using 20 \times fewer parameters (1.5B vs. 32B), with particularly strong performance on competition-level problems (AIME24, AIME25, AMC23) and MATH500. PieceHint-Qwen3-1.7B and PieceHint-DeepSeek-1.5B also demonstrate substantial improvements over their respective base models, outperforming significantly larger 4B baselines. By strategically identifying critical reasoning bottlenecks through value-driven piece selection, allocating hints based on problem difficulty, and progressively withdrawing scaffolding during training, small models can develop robust reasoning capabilities that rival those achieved through conventional parameter scaling.

5.3 Ablation Studies

Table 2: Comprehensive ablation studies on PieceHint components. All experiments use Nemotron-1.5B as the base model. The best result in each group is **bolded** with distinct colors.

Component	Method	AIME24	AIME25	AMC23	MATH500	Minerva	Olympiad	Avg.
<i>Problem Selection Strategy</i>								
Problem Selection	Hard Problems Only	49.8	39.2	84.6	88.1	26.3	63.5	58.6
	Random Selection	52.9	42.1	87.8	90.2	27.8	66.8	61.3
	Variance-Based	54.5	43.7	89.1	91.3	28.5	68.3	62.6
<i>Piece Selection Strategy</i>								
Piece Selection	50% Prefix Truncation	46.8	36.2	82.5	86.3	25.8	61.4	56.5
	Random Piece Selection	49.2	38.9	85.1	88.2	26.7	63.8	58.7
	25% Prefix Truncation	51.7	41.3	86.8	89.5	27.3	65.9	60.4
	Value-Driven	54.5	43.7	89.1	91.3	28.5	68.3	62.6
<i>Progressive Curriculum</i>								
Progressive Curriculum	w/o Withdrawal	52.1	41.5	87.3	89.8	27.6	66.2	60.8
	w/ Withdrawal	54.5	43.7	89.1	91.3	28.5	68.3	62.6

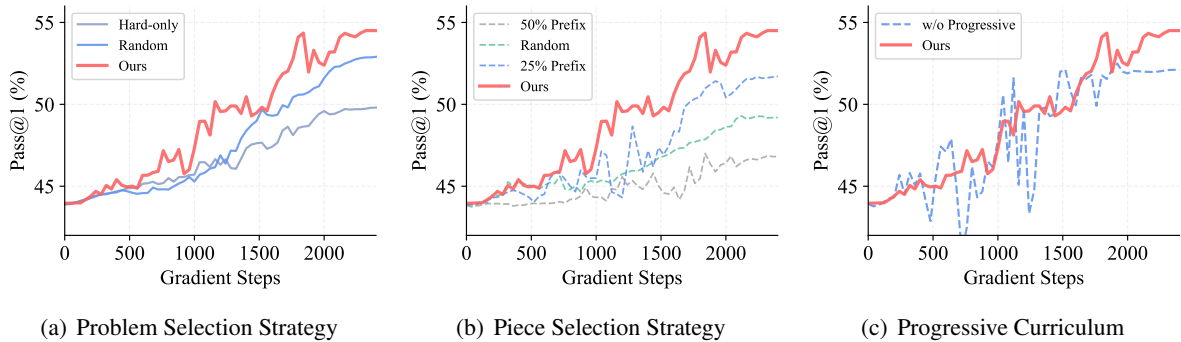


Figure 3: Training dynamics on AIME24 across ablation variants. All curves start from similar initial performance and diverge as training progresses, with PieceHint (red solid lines) consistently achieving the highest final performance while maintaining stable convergence.

We conduct ablation studies to validate the effectiveness of each component in PieceHint. We systematically evaluate three core design choices: variance-based problem selection, value-driven piece selection, and progressive curriculum. All ablation experiments use Nemotron-1.5B as the base model to ensure controlled comparison. Table 2 presents quantitative results, and Figure 3 illustrates the training dynamics.

Impact of Problem Selection Strategy. Table 2 (top) and Figure 3(a) compare three strategies. Hard-only includes intractable problems providing insufficient signals; random sampling achieves moderate results but lacks systematic difficulty targeting. Variance-based selection outperforms both by focusing on challenging yet learnable problems within the model’s capability range, validating that capability-aligned filtering is essential.

Impact of Value-Driven Piece Selection. Table 2 (middle) and Figure 3(b) compare value-driven piece identification against position-based and random selection baselines. Random selection samples pieces without considering their importance, while prefix truncation strategies select pieces based on sequential position. Among position-based methods, 50% prefix truncation performs worst by including excessive redundant early steps, while the 25% variant achieves better results but both miss critical bottlenecks appearing later. Random selection achieves moderate performance but lacks systematic targeting. Value-driven selection consistently outperforms all baselines by strategically identifying pivotal lemmas and key insights regardless of position, demonstrating that bottleneck-focused augmentation is substantially more effective than uniform coverage.

Impact of Progressive Curriculum. Table 2 (bottom section) and Figure 3(c) validate the ef-

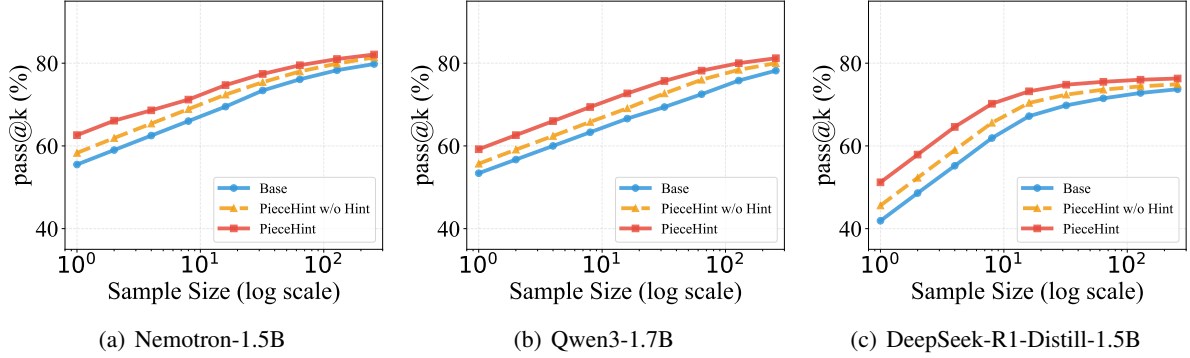


Figure 4: Pass@k performance across different sampling budgets showing convergent behavior at higher k values.

fectiveness of progressive piece withdrawal. The variant without progressive curriculum maintains fixed hint sets throughout training, leading to over-reliance on scaffolding and degraded performance when evaluated without hints. Progressive withdrawal achieves superior performance by gradually reducing hint dependency while preserving strategic guidance, enabling the model to internalize critical reasoning patterns and generalize robustly to hint-free scenarios.

5.4 Pass@k Analysis

Figure 4 shows pass@k performance across sampling budgets. PieceHint consistently outperforms both base models and standard RL training across all k values, with particularly pronounced gains at lower k values. The consistent improvement across increasing k demonstrates that strategic piece injection enhances solution diversity rather than creating hint dependency—by providing learning signals at critical bottlenecks and then progressively withdrawing guidance, models develop diverse reasoning strategies through independent exploration. As k increases, the performance gap gradually narrows, though PieceHint maintains advantages throughout the sampling range.

5.5 Case Study

Figure 5 analyzes hint strategies on a combinatorial problem (Appendix J). When models fail at critical reasoning bottlenecks, sparse rewards provide no learning signals. We compare: 50% prefix wastes hints on routine steps the model can handle; random selection hits bottlenecks inefficiently; PieceHint identifies where models actually struggle and strategically bridges these gaps. This demonstrates that effective augmentation requires understanding which specific reasoning barriers block learning,

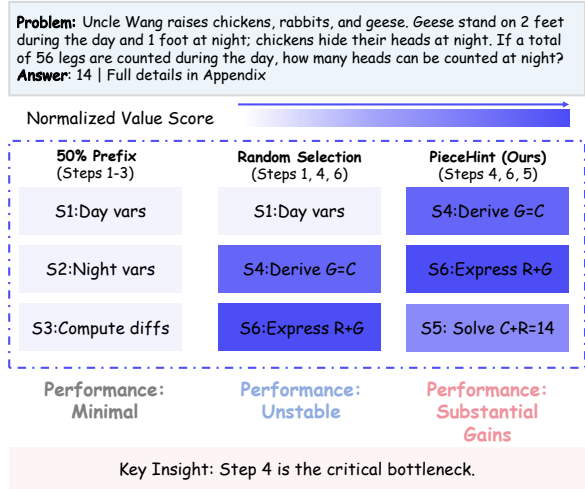


Figure 5: Hint strategy comparison (Appendix J). PieceHint targets actual learning bottlenecks, unlocking reward signals.

not just providing arbitrary solution fragments.

6 Conclusion

We present PieceHint, a value-driven framework that strategically provides hints at critical reasoning bottlenecks during reinforcement learning. Unlike uniform hint provision, PieceHint identifies where models actually struggle and bridges these gaps through value-driven piece selection, adaptive allocation, and progressive withdrawal, enabling transition from scaffolded to hint-free reasoning. Experiments on six mathematical benchmarks show our 1.5B model achieves comparable performance to 32B baselines while preserving solution diversity, demonstrating remarkable parameter efficiency. We further validate scalability to 7B models (Appendix F) and generalization to code generation (Appendix G), confirming that the core mechanism transfers across model scales and

reasoning domains. Future work includes investigating fully adaptive withdrawal strategies that dynamically adjust based on real-time model capability assessments.

Limitations

PieceHint demonstrates strong performance across 1.5B and 7B parameter models (Appendix F), and generalizes to code generation tasks beyond mathematics (Appendix G). Our frequency-based withdrawal uses a fixed schedule (N_{check}), trading adaptivity for simplicity. While adaptive strategies that dynamically adjust withdrawal rates and hint allocations based on real-time model capability could potentially improve performance, they would require continuous evaluation of model mastery on each problem to determine optimal training data and hint provision, significantly increasing engineering complexity and computational cost. Our fixed-schedule approach represents a practical trade-off between performance optimization and training efficiency. Additionally, PieceHint’s offline preprocessing relies on a stronger LLM as an external oracle; as shown in Appendix D, the framework is robust to the choice of scoring model, but this dependency represents a structural boundary condition of the method.

References

- Xin-Qiang Cai and Masashi Sugiyama. 2026. Vi-curl: Stabilizing verifier-independent rl reasoning via confidence-guided variance reduction. *arXiv preprint arXiv:2602.12579*.
- Xin-Qiang Cai, Wei Wang, Feng Liu, Tongliang Liu, Gang Niu, and Masashi Sugiyama. 2025. Reinforcement learning with verifiable yet noisy rewards under imperfect verifiers. *arXiv preprint arXiv:2510.00915*.
- Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde De Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, and 1 others. 2021. Evaluating large language models trained on code. *arXiv preprint arXiv:2107.03374*.
- Ruiyi Ding, Yongxuan Lv, Xianhui Meng, Jiahe Song, Chao Wang, Chen Jiang, and Yuan Cheng. 2026. Prpo: Aligning process reward with outcome reward in policy optimization. *arXiv preprint arXiv:2601.07182*.
- Tianqing Fang, Zhisong Zhang, Xiaoyang Wang, Rui Wang, Can Qin, Yuxuan Wan, Jun-Yu Ma, Ce Zhang, Jiaqi Chen, Xiyun Li, Hongming Zhang, Haitao Mi, and Dong Yu. 2025. Cognitive kernel-pro: A framework for deep research agents and agent foundation models training. *CoRR*, abs/2508.00414.
- Yangyi Fang, Jiaye Lin, Xiaoliang Fu, Cong Qin, Haolin Shi, Chaowen Hu, Lu Pan, Ke Zeng, and Xunliang Cai. 2026a. How to allocate, how to learn? dynamic rollout allocation and advantage modulation for policy optimization. *arXiv preprint arXiv:2602.19208*.
- Yangyi Fang, Jiaye Lin, Xiaoliang Fu, Cong Qin, Haolin Shi, Chang Liu, and Peilin Zhao. 2026b. Proximity-based multi-turn optimization: Practical credit assignment for llm agent training. *arXiv preprint arXiv:2602.19225*.
- Xiaoliang Fu, Jiaye Lin, Yangyi Fang, Chaowen Hu, Cong Qin, Zekai Shao, Binbin Zheng, Lu Pan, and Ke Zeng. 2026a. From $\log \pi$ to π : Taming divergence in soft clipping via bilateral decoupled decay of probability gradient weight. *arXiv preprint arXiv:2603.14389*.
- Xiaoliang Fu, Jiaye Lin, Yangyi Fang, Binbin Zheng, Chaowen Hu, Zekai Shao, Cong Qin, Lu Pan, Ke Zeng, and Xunliang Cai. 2026b. Maspo: Unifying gradient utilization, probability mass, and signal reliability for robust and sample-efficient llm reasoning. *arXiv preprint arXiv:2602.17550*.
- Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, and 1 others. 2025a. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*.
- Ruohao Guo, Afshin Oroojlooy, Roshan Sridhar, Miguel Ballesteros, Alan Ritter, and Dan Roth. 2025b. Tree-based dialogue reinforced policy optimization for red-teaming attacks. *arXiv preprint arXiv:2510.02286*.
- Chaoqun He, Renjie Luo, Yuzhuo Bai, Shengding Hu, Zhen Leng Thai, Junhao Shen, Jinyi Hu, Xu Han, Yujie Huang, Yuxiang Zhang, and 1 others. 2024. Olympiadbench: A challenging benchmark for promoting agi with olympiad-level bilingual multimodal scientific problems. *arXiv preprint arXiv:2402.14008*.
- Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. 2021. Measuring mathematical problem solving with the math dataset. *arXiv preprint arXiv:2103.03874*.
- Zhenyu Hou, Ziniu Hu, Yujiang Li, Rui Lu, Jie Tang, and Yuxiao Dong. 2025. Treerl: Llm reinforcement learning with on-policy tree search. *Preprint*, arXiv:2506.11902.
- Zixuan Huang, Yikun Ban, Lean Fu, Xiaojie Li, Zhongxiang Dai, Jianxin Li, and Deqing Wang. 2025. Adaptive sample scheduling for direct preference optimization. *arXiv preprint arXiv:2506.17252*.

- Zixuan Huang, Xin Xia, Yuxi Ren, Jianbin Zheng, Xuanda Wang, Zhixia Zhang, Hongyan Xie, Songshi Liang, Zehao Chen, Xuefeng Xiao, and 1 others. 2026a. Does your reasoning model implicitly know when to stop thinking? *arXiv preprint arXiv:2602.08354*.
- Zixuan Huang, Xin Xia, Yuxi Ren, Jianbin Zheng, Xuefeng Xiao, Hongyan Xie, Huaqiu Li, Songshi Liang, Zhongxiang Dai, Fuzhen Zhuang, Jianxin Li, Yikun Ban, and Deqing Wang. 2026b. [Real-time aligned reward model beyond semantics](#).
- HuggingFace. 2025. [Math-verify](#).
- Amirhossein Kazemnejad, Milad Aghajohari, Eva Portelance, Alessandro Sordani, Siva Reddy, Aaron Courville, and Nicolas Le Roux. 2025. [Vineppo: Refining credit assignment in rl training of llms](#). *Preprint*, arXiv:2410.01679.
- Aitor Lewkowycz, Anders Andreassen, David Dohan, Ethan Dyer, Henryk Michalewski, Vinay Ramasesh, Ambrose Slone, Cem Anil, Imanol Schlag, Theo Gutman-Solo, and 1 others. 2022. Solving quantitative reasoning problems with language models. *Conference on Neural Information Processing Systems (NeurIPS)*.
- Mengdi Li, Jiaye Lin, Xufeng Zhao, Wenhao Lu, Peilin Zhao, Stefan Wermtner, and Di Wang. 2025. Curriculum-rlaif: Curriculum alignment with reinforcement learning from ai feedback. *arXiv preprint arXiv:2505.20075*.
- Mukai Li, Qingcheng Zeng, Tianqing Fang, Zhenwen Liang, Linfeng Song, Qi Liu, Haitao Mi, and Dong Yu. 2026. [Verified critical step optimization for LLM agents](#). *CoRR*, abs/2602.03412.
- Dengcan Liu, Fengkai Yang, Xiaohan Wang, Shurui Yan, Jiajun Chai, Jiahao Li, Yikun Ban, Zhendong Mao, Wei Lin, and Guojun Yin. 2026. [Cdrmm: Contrast-driven rubric generation for reliable and interpretable reward modeling](#). *arXiv preprint arXiv:2603.08035*.
- Mingjie Liu, Shizhe Diao, Ximing Lu, Jian Hu, Xin Dong, Yejin Choi, Jan Kautz, and Yi Dong. 2025a. [Prorl: Prolonged reinforcement learning expands reasoning boundaries in large language models](#). *arXiv preprint arXiv:2505.24864*.
- Peiyang Liu. 2024. Unsupervised corrupt data detection for text training. *Expert Systems with Applications*, 248:123335.
- Peiyang Liu, Xi Wang, Ziqiang Cui, and Wei Ye. 2025b. [Queries are not alone: Clustering text embeddings for video search](#). In *Proceedings of the 48th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 874–883.
- Peiyang Liu, Jinyu Yang, Lin Wang, Sen Wang, Yunlai Hao, and Huihui Bai. 2023. Retrieval-based unsupervised noisy label detection on text data. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*, pages 4099–4104.
- Zichen Liu, Changyu Chen, Wenjun Li, Penghui Qi, Tianyu Pang, Chao Du, Wee Sun Lee, and Min Lin. 2025c. [Understanding r1-zero-like training: A critical perspective](#). *arXiv preprint arXiv:2503.20783*.
- Michael Luo, Sijun Tan, Justin Wong, Xiaoxiang Shi, William Y. Tang, Manan Roongta, Colin Cai, Jeffrey Luo, Li Erran Li, Raluca Ada Popa, and Ion Stoica. 2025. [Deepscaler: Surpassing o1-preview with a 1.5b model by scaling rl](#). Notion Blog.
- MAA. 2023. [American mathematics competitions - amc](#).
- MAA. 2025. [American invitational mathematics examination - aime](#).
- Saumya Malik, Valentina Pyatkin, Sander Land, Jacob Morrison, Noah A. Smith, Hannaneh Hajishirzi, and Nathan Lambert. 2025. [Rewardbench 2: Advancing reward model evaluation](#). *Preprint*, arXiv:2506.01937.
- Ivan Moshkov, Darragh Hanley, Ivan Sorokin, Shubham Toshniwal, Christof Henkel, Benedikt Schifferer, Wei Du, and Igor Gitman. 2025. [Aimo-2 winning solution: Building state-of-the-art mathematical reasoning models with openmathreasoning dataset](#). *arXiv preprint arXiv:2504.16891*.
- Shubham Parashar, Shurui Gui, Xiner Li, Hongyi Ling, Sushil Vemuri, Blake Olson, Eric Li, Yu Zhang, James Caverlee, Dileep Kalathil, and Shuiwang Ji. 2025. [Curriculum reinforcement learning from easy to hard tasks improves llm reasoning](#).
- Yuxiao Qu, Matthew Y. R. Yang, Amrith Setlur, Lewis Tunstall, Edward Emanuel Beeching, Ruslan Salakhutdinov, and Aviral Kumar. 2025. [Optimizing test-time compute via meta reinforcement fine-tuning](#). *Preprint*, arXiv:2503.07572.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. [Proximal policy optimization algorithms](#). *Preprint*, arXiv:1707.06347.
- Rulin Shao, Shuyue Stella Li, Rui Xin, Scott Geng, Yiping Wang, Sewoong Oh, Simon Shaolei Du, Nathan Lambert, Sewon Min, Ranjay Krishna, and 1 others. 2025. [Spurious rewards: Rethinking training signals in rlvr](#). *arXiv preprint arXiv:2506.10947*.
- Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, YK Li, Y Wu, and 1 others. 2024. [Deepseek-math: Pushing the limits of mathematical reasoning in open language models](#). *arXiv preprint arXiv:2402.03300*.

- Guangming Sheng, Chi Zhang, Zilingfeng Ye, Xibin Wu, Wang Zhang, Ru Zhang, Yanghua Peng, Haibin Lin, and Chuan Wu. 2025. Hybridflow: A flexible and efficient rlhf framework. In *European Conference on Computer Systems (EuroSys)*.
- Yuqiao Tan, Minzheng Wang, Shizhu He, Huanxuan Liao, Chengfeng Zhao, Qiunan Lu, Tian Liang, Jun Zhao, and Kang Liu. 2025. Bottom-up policy optimization: Your language model policy secretly contains internal policies. *arXiv preprint arXiv:2512.19673*.
- Peiyi Wang, Lei Li, Zhihong Shao, R. X. Xu, Damai Dai, Yifei Li, Deli Chen, Y. Wu, and Zhifang Sui. 2024. Math-shepherd: Verify and reinforce llms step-by-step without human annotations. *Preprint, arXiv:2312.08935*.
- Xumeng Wen, Zihan Liu, Shun Zheng, Zhijian Xu, Shengyu Ye, Zhirong Wu, Xiao Liang, Yang Wang, Junjie Li, Ziming Miao, and 1 others. 2025. Reinforcement learning with verifiable rewards implicitly incentivizes correct reasoning in base llms. *arXiv preprint arXiv:2506.14245*.
- Jianhao Yan, Yafu Li, Zican Hu, Zhi Wang, Ganqu Cui, Xiaoye Qu, Yu Cheng, and Yue Zhang. 2025a. Learning to reason under off-policy guidance. *arXiv preprint arXiv:2504.14945*.
- Yuchen Yan, Yongliang Shen, Yang Liu, Jin Jiang, Mengdi Zhang, Jian Shao, and Yueting Zhuang. 2025b. Infythink: Breaking the length limits of long-context reasoning in large language models. *arXiv preprint arXiv:2503.06692*.
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, and 1 others. 2025. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*.
- Fengkai Yang, Zherui Chen, Xiaohan Wang, Xiaodong Lu, Jiajun Chai, Guojun Yin, Wei Lin, Shuai Ma, Fuzhen Zhuang, Deqing Wang, and 1 others. 2026. Your group-relative advantage is biased. *arXiv preprint arXiv:2601.08521*.
- Qiyang Yu, Zheng Zhang, Ruofei Zhu, Yufeng Yuan, Xiaochen Zuo, Yu Yue, Weinan Dai, Tiantian Fan, Gao-hong Liu, Lingjun Liu, and 1 others. 2025a. Dapo: An open-source llm reinforcement learning system at scale. *arXiv preprint arXiv:2503.14476*.
- Wenhao Yu, Zhenwen Liang, Chengsong Huang, Kishan Panaganti, Tianqing Fang, Haitao Mi, and Dong Yu. 2025b. Guided self-evolving llms with minimal human supervision. *CoRR, abs/2512.02472*.
- Zhiqi Yu, Zhangquan Chen, Mengting Liu, Heye Zhang, and Liangqiong Qu. 2026. Unveiling implicit advantage symmetry: Why grpo struggles with exploration and difficulty adaptation. *arXiv preprint arXiv:2602.05548*.
- Yang Yue, Zhiqi Chen, Rui Lu, Andrew Zhao, Zhaokai Wang, Shiji Song, and Gao Huang. 2025a. Does reinforcement learning really incentivize reasoning capacity in llms beyond the base model? *arXiv preprint arXiv:2504.13837*.
- Yang Yue, Zhiqi Chen, Rui Lu, Andrew Zhao, Zhaokai Wang, Shiji Song, and Gao Huang. 2025b. Does reinforcement learning really incentivize reasoning capacity in llms beyond the base model? *arXiv preprint arXiv:2504.13837*.
- Kaiyi Zhang, Ang Lv, Jinpeng Li, Yongbo Wang, Feng Wang, Haoyuan Hu, and Rui Yan. 2025. Stephint: Multi-level stepwise hints enhance reinforcement learning to reason. *arXiv preprint arXiv:2507.02841*. Submitted Jul 3, 2025.
- Zhixia Zhang, Zixuan Huang, Xin Xia, Deqing Wang, Fuzhen Zhuang, Shuai Ma, Ning Ding, Yaodong Yang, Jianxin Li, and Yikun Ban. 2026. Heterogeneous agent collaborative reinforcement learning. *arXiv preprint arXiv:2603.02604*.
- Rosie Zhao, Alexandru Meterez, Sham Kakade, Cengiz Pehlevan, Samy Jelassi, and Eran Malach. 2025. Echo chamber: RL post-training amplifies behaviors learned in pretraining.
- Yixiao Zhou, Yang Li, Dongzhou Cheng, Hehe Fan, and Yu Cheng. 2026. Look inward to explore outward: Learning temperature policy from llm internal states via hierarchical rl. *arXiv preprint arXiv:2602.13035*.
- Yixiao Zhou, Ziyu Zhao, Dongzhou Cheng, Zhiliang Wu, Jie Gui, Yi Yang, Fei Wu, Yu Cheng, and Hehe Fan. 2025. Dropping experts, recombining neurons: Retraining-free pruning for sparse mixture-of-experts llms. In *Findings of the Association for Computational Linguistics: EMNLP 2025*, pages 15169–15186.
- Xinyu Zhu, Mengzhou Xia, Zhepei Wei, Wei-Lin Chen, Danqi Chen, and Yu Meng. 2025. The surprising effectiveness of negative reinforcement in llm reasoning. *arXiv preprint arXiv:2506.01347*.

A Additional Related Work

This section surveys concurrent and complementary work in reinforcement learning for LLM reasoning that further contextualizes PieceHint.

Policy optimization and gradient stability. A rich body of recent work improves the optimization dynamics of RL-based LLM training. Several studies identify structural issues in standard advantage estimation: Yang et al. (2026) analyze biases inherent in group-relative advantage computation, while Yu et al. (2026) reveal implicit symmetry problems in GRPO that hinder exploration and difficulty adaptation—a challenge that PieceHint addresses from the input side by providing targeted scaffolding

rather than relying solely on the optimizer. At the gradient level, [Fu et al. \(2026a\)](#) address divergence in soft clipping objectives via bilateral decoupled decay of probability gradient weights, and [Fu et al. \(2026b\)](#) unify gradient utilization, probability mass management, and signal reliability for more robust reasoning. Complementary to these optimization-level methods, [Tan et al. \(2025\)](#) propose bottom-up policy optimization that leverages internal policies latent within pre-trained language models. Dynamic resource allocation has also been explored: [Fang et al. \(2026a\)](#) study how to jointly optimize rollout allocation and advantage modulation, which is orthogonal to PieceHint’s input-level hint allocation and could be combined to further improve sample efficiency. These approaches all operate at the optimization level and are compatible with PieceHint as drop-in enhancements to the training pipeline.

Reward modeling and credit assignment. Reliable reward signals are critical for RL training. [Cai et al. \(2025\)](#) study RL under verifiable yet noisy rewards from imperfect verifiers, and [Cai and Sugiyama \(2026\)](#) propose confidence-guided variance reduction to stabilize verifier-independent RL reasoning—both addressing the signal quality issues that motivate PieceHint’s use of binary outcome rewards with strategic hint injection. [Ding et al. \(2026\)](#) propose aligning process-level and outcome-level rewards in policy optimization, offering a supervision signal complementary to PieceHint’s piece-level scaffolding. [Liu et al. \(2026\)](#) develop contrast-driven rubric generation for more reliable and interpretable reward modeling, and [Huang et al. \(2026a\)](#) investigate whether reasoning models can implicitly learn to regulate their own thinking depth—both relevant to understanding the interaction between hint provision and autonomous reasoning development. Beyond single-turn settings, [Fang et al. \(2026b\)](#) address credit assignment in multi-turn LLM agent training, sharing the core motivation with PieceHint that sparse terminal rewards are insufficient for effective learning. On the preference alignment side, [Huang et al. \(2025\)](#) study adaptive sample scheduling for direct preference optimization, and [Huang et al. \(2026b\)](#) explore reward modeling beyond semantic alignment. **Exploration, temperature control, and model efficiency.** Effective exploration is essential when reward signals are sparse. [Zhou et al. \(2026\)](#) propose learning temperature policies from LLM internal states via hierarchical RL, dynamically adapting

sampling diversity in a manner complementary to PieceHint’s progressive hint withdrawal. [Zhang et al. \(2026\)](#) study collaborative RL among heterogeneous agents, extending the multi-agent perspective to settings where diverse reasoning strategies must be coordinated. On model efficiency, [Zhou et al. \(2025\)](#) propose retraining-free pruning for sparse mixture-of-experts models, enabling deployment of capable base models at reduced cost. Finally, [Guo et al. \(2025b\)](#) apply tree-structured dialogue policy optimization to red-teaming, demonstrating the breadth of domains where RL-based policy learning continues to advance.

B Training Configuration Details

Tables 3 and 4 present the complete training configurations. Table 3 shows general hyperparameters shared across all experiments, while Table 4 presents model-specific problem selection and piece allocation parameters.

C Direct Comparison with StepHint

To isolate PieceHint’s gains from the choice of hint strategy, we conduct a controlled head-to-head comparison with StepHint ([Zhang et al., 2025](#)) under identical experimental conditions: the same base model (Nemotron-1.5B), training dataset, hyperparameters, and computational budget. StepHint prepends the first 25% of ground-truth solution steps as positional prefix hints—the best-performing variant of StepHint as identified in our ablation studies (Table 2).

Table 5 shows that PieceHint consistently outperforms StepHint across all six benchmarks, with an average improvement of 3.8 points (62.6 vs. 58.8). The gap is most pronounced on competition-level benchmarks (AIME24: +4.7, AIME25: +4.2, AMC23: +3.7), where critical bottlenecks tend to appear later in solutions and are systematically missed by positional prefix selection. These results confirm that value-driven identification of reasoning bottlenecks provides a clear advantage over position-based hint allocation: mechanically truncating the first $k\%$ of steps may cover routine early derivations while leaving the pivotal insights—which often appear mid-solution—unscaffolded.

D Scoring Robustness Analysis

PieceHint relies on an LLM-based scoring step to assign value scores to reasoning pieces. We analyze

two potential sources of instability: the choice of scoring model and sensitivity to prompt phrasing.

Cross-Model Scoring Consistency

We preprocess 1,000 training problems using three different LLM backends and measure both end-to-end performance and the agreement among top- k piece selections.

Performance variance across LLM backends is minimal (≤ 0.6 average points), confirming that the framework is robust to scoring model choice. To understand why, we further measure cross-scorer agreement on 500 problems:

High piece-selection overlap (82–87%) and strong rank correlations ($\rho = 0.80$ – 0.86) indicate that different LLMs converge on similar high-value bottlenecks. The framework’s resilience stems from top- k selection: as long as scorers agree on which pieces are high-value—which they consistently do—minor disagreements on lower-ranked pieces do not affect training. We additionally ran GPT-4 preprocessing three times on the same 500 problems at temperature=0, achieving 96% top-3 overlap across runs (std = 3.3%, Pearson $r = 0.94 \pm 0.02$), confirming strong within-model reproducibility despite minor API non-determinism.

Prompt Sensitivity Analysis

We test five prompt variants that change wording while preserving the same three scoring criteria (novelty, difficulty, impact). The *Piece Overlap Rate* measures the fraction of problems where the top-3 selected pieces match the original prompt’s selections.

Despite variations in prompt phrasing, piece selection remains highly consistent (85–89% overlap) and final performance stays stable with ≤ 0.6 point variation. This demonstrates that PieceHint is robust to prompt phrasing as long as the core scoring criteria are preserved.

Preprocessing cost and failure modes. Scoring is a one-time offline step requiring one LLM API call per training problem; results are fully cacheable for all subsequent training runs. We analyzed failure modes on 100 randomly sampled problems and found that the most common failure is overgranular decomposition, where a single insight is split into multiple pieces. Our min-max normalization and top- k selection handle such granularity differences and noisy scores effectively.

E Robustness to Corrupted Hints

A critical question is whether incorrect or adversarial pieces can cause catastrophic failure and whether RL training might amplify early hint bias. We address both concerns directly.

Performance Under Corrupted Hints

We deliberately inject errors into piece decomposition and scoring to simulate realistic failure modes, using Nemotron-1.5B as the base model.

Even under adversarial conditions where hints actively mislead, performance drops remain modest (at most -3.9 on AIME24 and -2.7 on MATH500). Two protective mechanisms explain this graceful degradation: (1) RL exploration discovers correct solutions despite misleading hints via binary outcome rewards that signal when hint-following fails, and (2) the scaffolding regime provides intermediate checkpoints rather than end-to-end solutions, requiring models to bridge gaps through independent reasoning and thereby limiting the damage any single incorrect piece can cause.

Progressive Withdrawal Prevents Bias Amplification

If RL were amplifying hint dependency, we would expect models trained with fixed hints to perform *worse* when hints are removed at evaluation. The ablation in Table 2 already demonstrates the opposite: progressive withdrawal (62.6 avg.) substantially outperforms fixed hints (60.8 avg.), confirming that models learn generalizable reasoning strategies rather than hint-specific shortcuts.

Further evidence comes from our $\text{pass}@k$ analysis (Section 5.4): if hints biased exploration, $\text{pass}@k$ curves would degrade at higher k as models converge to hint-dependent trajectories. Instead, we observe sustained gains across all k values (Figure 4), confirming enhanced solution diversity.

The mechanism preventing bias amplification involves three factors: (1) binary outcome rewards correct hint-following failures, (2) RL exploration discovers alternative solution paths beyond scaffolded trajectories, and (3) gradual withdrawal forces trajectory revalidation, preventing memorization of hint-specific patterns.

F Scalability to Larger Models

While our main experiments focus on 1.5B parameter models where targeted hint allocation provides

the most benefit, we validate PieceHint’s scalability to larger models. We apply the same PieceHint pipeline to DeepSeek-R1-Distill-7B, using its own checkpoint as the capability reference model μ_0 for variance-based selection.

PieceHint achieves the highest performance across all benchmarks, outperforming both the backbone and GRPO baseline. The framework delivers substantial improvements on competition-level benchmarks (AIME24: +14.7, AIME25: +14.0, AMC23: +14.6 over backbone) and meaningful gains on reasoning-focused tasks (Minerva: +3.4, Olympiad: +4.4).

As expected, relative gains are somewhat smaller than for 1.5B models: stronger 7B models encounter fewer critical bottlenecks per problem, so targeted hint allocation provides proportionally less benefit. The capability-aligned problem selection automatically accommodates this by assigning fewer hints to more capable models—validating that the framework scales gracefully across model sizes without manual recalibration.

G Code Generation Experiments

To validate that the core insight of PieceHint generalizes beyond mathematical reasoning, we apply the framework to code generation. We use DeepSeek-R1-Distill-Qwen-7B as the base model and evaluate on LiveCodeBench and CodeForces.

Domain adaptation. PieceHint transfers naturally to code: variance-based problem selection identifies appropriately challenging coding tasks, value-driven piece identification scores algorithmic steps by complexity and novelty (e.g., recognizing the appropriate data structure, identifying an edge case invariant, applying a non-obvious algorithm), and progressive withdrawal prevents over-reliance on code scaffolding. Domain adaptation requires only adjusting piece decomposition and scoring criteria to code-specific characteristics.

PieceHint achieves the highest performance across both benchmarks, substantially outperforming the GRPO baseline (LCB Avg@16: +4.9, LCB Pass@16: +3.7, CF Rating: +238, CF Percentile: +16.5 pp). These results confirm that the core mechanism—strategically scaffolding critical reasoning bottlenecks then progressively withdrawing scaffolding—generalizes to structured problem-solving domains beyond mathematics, provided the domain supports verifiable outcome rewards and decomposable solution steps.

H Software Dependencies

Table 12 lists the key software packages and their versions used in our implementation.

Our implementation is available at [anonymized for review] and will be released upon publication with a complete requirements.txt and Docker environment for reproducibility.

I Value Scoring Prompt

We employ GPT-4 to assign discrete value scores $V(p_i) \in \{1, 2, 3, 4, 5\}$ to each reasoning piece by holistically considering three criteria with emphasis on impact. Raw scores are then normalized within each problem using min-max normalization (Eq. 6).

Value Scoring Prompt Template

Rate the importance of this reasoning step on a scale of 1-5.
Consider three criteria (prioritize Impact):

- **Novelty:** Is this insight non-trivial or creative?
- **Difficulty:** Does it require mathematical maturity?
- **Impact:** Do subsequent steps critically depend on it?

Scoring guidelines:

- 1 = Trivial (e.g., restating the answer)
- 2 = Routine (standard manipulations)
- 3 = Moderate (requires reasoning, not critical)
- 4 = Important (valuable insight)
- 5 = Critical bottleneck (blocks solution if missing)

Problem: *[Insert problem statement]*

Reasoning step: *[Insert piece text]*

Output: A single integer from 1 to 5

Example Application:

Problem: Some of 100 towns are connected by roads. For each two towns A and B connected by a road, there exists a town C not connected to at least one of A or B . What is the maximum possible number of roads?

Decomposed Steps & Scores:

1. Restate condition in graph theory terms. **Score: 2**
2. Recognize this prevents complete graph. **Score: 2**
3. Hypothesize removing perfect matching. **Score: 4**
4. Calculate: $\binom{100}{2} - 50 = 4900$ edges. **Score: 2**
5. Verify construction satisfies condition. **Score: 5**
6. Conclude maximum is 4900. **Score: 1**

Normalized: $V_{\min} = 1, V_{\max} = 5 \Rightarrow D(p_i): [0.25, 0.25, 0.75, 0.25, 1.00, 0.00]$
 Step 5 receives the highest score as it critically establishes correctness—without verification, the construction remains unproven.

J Case Study Problem Details

This appendix provides complete details for the case study analyzed in the main text. We present a combinatorial reasoning problem where models must derive a non-obvious invariance relationship. Table 13 shows the 7 reasoning steps with normalized value scores. Table 14 compares three hint allocation strategies, demonstrating that value-driven selection yields $7\times$ better outcomes than position-based coverage.

Problem Statement

Context: Uncle Wang raises chickens (C), rabbits (R), and geese (G). Geese stand on both feet during the day, one foot at night. Chickens hide their heads when sleeping. Rabbits show no behavioral changes.

Observation: The difference between feet and heads remains constant day and night.

Given: 56 legs counted during the day.

Question: How many heads at night?

Answer: 14

Solution Process (7 Steps)

Let C, R, G denote chickens, rabbits, geese.

Steps 1-2: Express day (heads: $C + R + G$, feet: $2C + 4R + 2G$, diff: $C + 3R + G$) and night (heads: $R + G$, feet: $2C + 4R + G$, diff: $2C + 3R$).

Step 3: Day diff = $C + 3R + G$; Night diff = $2C + 3R$.

Step 4 (Critical): Since equal: $C + 3R + G = 2C + 3R \Rightarrow \boxed{G = C}$

Step 5: From $2C + 4R + 2G = 56$ and $G = C$: $4C + 4R = 56 \Rightarrow C + R = 14$

Steps 6-7: Night heads = $R + G = R + C = 14$

Why Step 4 Gets Highest Score ($D = 0.94$)

Novelty: Non-obvious that "same difference" implies $G = C$; requires abstract algebraic reasoning.

Difficulty: Must manipulate two complex equations simultaneously; error-prone.

Impact: Gateway step—without $G = C$, all subsequent steps blocked.

Empirical: Without Step 4: 0/16 succeed (0%); with Step 4: 7/16 succeed (43.8%). Step 4 alone yields 28.1%, vs. Steps 1-3 combined (6.2%).

Strategy Comparison (3 hints each)

(1) Position-based fails: 50% prefix misses critical Step 4 (6.2% pass@16).

(2) Random unstable: Includes Step 4 with only 33% probability.

(3) Value-driven succeeds: Always provides Step 4, achieving 43.8% ($7.1\times$ improvement).

Conclusion: Strategic allocation at high-value bottlenecks yields $7\times$ better outcomes.

Table 3: General training hyperparameters and system configuration (shared across all models).

PieceHint-Specific Parameters			
Withdrawal Checkpoint (N_{check})	2	Max Initial Pieces (k_{max})	3
Model Configuration			
Gradient Checkpointing	Enabled	Activation Offload	Disabled
LoRA Rank	0 (Disabled)	Fused Kernels	Disabled
Remove Padding	Enabled	FSDP Strategy	Fully Sharded
Forward Prefetch	Disabled	Reshard After Forward	Enabled
Parameter Offload	Disabled		
Optimization			
Learning Rate	1×10^{-6}	Weight Decay	0.1
Gradient Clipping	1.0	LR Warmup Steps	10
Warmup Style	Constant	Min LR Ratio	0.0
GRPO Configuration			
Mini Batch Size	32	Clip Ratio	0.2
KL Loss Coefficient	0.0	Advantage Estimator	GRPO
Rollout & Generation			
Inference Engine	vLLM	Tensor Parallel Size	2
Samples per Prompt (n)	16	Temperature (τ)	1.0
Top-p (p)	1.0	Top-k	-1 (Disabled)
Max Prompt Length	2048	Max Response Length	8192
Do Sample	True	Ignore EOS	False
GPU Memory Utilization	0.8	Max Batched Tokens	10240
Max Num Sequences	1024	Chunked Prefill	Enabled
Enforce Eager Mode	True	Load Format	dummy_dtensor
Data Configuration			
Train Batch Size	512	Validation Batch Size	512
Dataloader Workers	8	Shuffle Training Data	True
Max Prompt Length	2048 tokens	Max Response Length	8192 tokens
Hardware & Training Schedule			
Number of Nodes	1	GPUs per Node	8
GPU Type	A100 80GB	Total GPUs	8
Total Gradient Steps	2,400	Total Rollout Samples	1M

Table 4: Model-specific problem selection and piece allocation parameters. Each model uses its own checkpoint as the capability reference model μ_0 for variance-based selection.

Parameter	PieceHint-DeepSeek-1.5B	PieceHint-Qwen3-1.7B	PieceHint-Nemotron-1.5B
<i>Problem Selection Thresholds (Section 4.1)</i>			
α_1 (hard filter)	0.2	0.2	0.2
α_2 (min capability)	0.1	0.1	0.1
α_3 (max capability)	0.4	0.4	0.4
<i>Piece Allocation Thresholds (Section 4.3)</i>			
β_1 (hard tier boundary)	0.15	0.15	0.15
β_2 (medium tier boundary)	0.35	0.35	0.35

Table 5: Controlled comparison between PieceHint and StepHint under identical settings. Base model: Nemotron-1.5B.

Method	AIME24	AIME25	AMC23	MATH500	Minerva	Olympiad	Avg.
No Hints (baseline)	44.0	34.2	82.7	88.5	22.3	61.2	55.5
StepHint (25% prefix)	49.8	39.5	85.4	88.9	25.1	64.2	58.8
PieceHint (ours)	54.5	43.7	89.1	91.3	28.5	68.3	62.6

Table 6: Performance across different LLM scoring backends (Base model: Nemotron-1.5B).

Scoring Model	AIME24	AIME25	AMC23	MATH500	Minerva	Olympiad	Avg.
GPT-4 (original)	54.5	43.7	89.1	91.3	28.5	68.3	62.6
DeepSeek-V3.2	54.1	43.3	88.7	91.0	28.2	67.9	62.2
Gemini-3-Flash-Preview	53.8	43.0	88.5	90.8	28.1	67.7	62.0

Table 7: Cross-LLM scoring consistency on 500 problems.

Scorer Pair	Top-3 Piece Overlap	Spearman ρ
GPT-4 \leftrightarrow DeepSeek-V3.2	87%	0.86
GPT-4 \leftrightarrow Gemini-3-Flash	84%	0.83
DeepSeek \leftrightarrow Gemini	82%	0.80

Table 8: Robustness to prompt variations (Base model: Nemotron-1.5B).

Prompt Variant	AIME24	MATH500	Piece Overlap
Original	54.5	91.3	100% (baseline)
Variant A (simpler language)	54.1	91.0	87%
Variant B (more detailed criteria)	54.3	91.2	89%
Variant C (different ordering)	53.9	90.9	85%
Variant D (examples-based)	54.2	91.1	88%

Table 9: Performance under deliberately corrupted hints (Base model: Nemotron-1.5B).

Corruption Type	AIME24	MATH500
Clean hints	54.5	91.3
25% wrong boundaries	52.8	90.1
Randomized scores	50.8	88.6
Worst pieces selected	51.2	89.4
50% corrupted hints	52.8	90.3
Contradictory hints	50.6	88.9

Table 10: Scalability experiments on 7B models. Base model: DeepSeek-R1-Distill-7B.

Method	AIME24	AIME25	AMC23	MATH500	Minerva	Olympiad	Avg.
DeepSeek-R1-Distill-7B (backbone)	39.6	29.9	78.8	80.3	36.8	55.1	53.4
GRPO	48.2	37.4	88.1	83.8	38.6	56.8	58.8
PieceHint (ours)	54.3	43.9	93.4	86.2	40.2	59.5	62.9

Table 11: Code generation results on LiveCodeBench and CodeForces. Base model: DeepSeek-R1-Distill-Qwen-7B.

Method	LCB Avg@16	LCB Pass@16	CF Rating	CF Percentile
Backbone (DeepSeek-R1-Distill-Qwen-7B)	28.4	47.2	892	18.7%
GRPO	31.9	51.6	1285	65.4%
PieceHint (ours)	36.8	55.3	1523	81.9%

Table 12: Software dependencies and environment configuration.

Core Framework			
Python	3.10+	PyTorch	2.5.0
Transformers	4.46+	Accelerate	1.2.1
PEFT	0.14+	Flash Attention	2.7.3
Distributed Training & RL			
VeRL	latest	Ray	2.40+
TensorDict	0.8.0–0.9.1	NVIDIA NCCL	cu12
vLLM	latest	Liger Kernel	latest
Data Processing			
Datasets (HF)	3.2+	PyArrow	$\geq 19.0.0$
Pandas	latest	Dill	latest
NumPy	$< 2.0.0$	TorchData	latest
Mathematical Verification			
Math-Verify	latest	Sympy	latest
Latex2Sympy2 (Extended)	latest	PyLaTeXEnc	latest
Word2Number	latest		
Logging & Experiment Management			
Weights & Biases	latest	TensorBoard	$\geq 2.19.0$
Hydra Core	latest	CodeTiming	latest
API & Utilities			
OpenAI	latest	FastAPI	latest
Uvicorn	latest	Packaging	≥ 20.0
Pre-commit	latest	Pybind11	latest

Step	Component	$D(p_i)$	Type
1	Express day scenario	0.25	Routine
2	Express night scenario	0.25	Routine
3	Compute differences	0.25	Routine
4	Derive $G = C$	0.94	Critical
5	Solve for $C + R$	0.56	Medium
6	Express $R + G$	0.75	Important
7	State answer	0.00	Trivial

Table 13: Decomposed steps with value scores.

Strategy	Hints	Avg D	Pass@16	Gain
No Hints	—	—	0.0%	—
50% Prefix	1,2,3	0.25	6.2%	1.0×
Random	1,4,6	0.65	18.8%	3.0×
25% Prefix	1,2	0.25	3.1%	0.5×
PieceHint	4,6,5	0.75	43.8%	7.1×

Table 14: Hint strategy comparison.