

LEPO: Latent Reasoning Policy Optimization for Large Language Models

Yuyan Zhou^{1,†} Jiarui Yu^{1,†} Hande Dong^{1*} Zhezheng Hao^{1,2}

Hong Wang¹ Jianqing Zhang³ Qiang Lin¹

¹ Tencent ² Zhejiang University

³ Shanghai Jiao Tong University

Abstract

Recently, latent reasoning has been introduced into large language models (LLMs) to leverage rich information within a continuous space. However, without stochastic sampling, these methods inevitably collapse to deterministic inference, failing to discover diverse reasoning paths. To bridge the gap, we inject controllable stochasticity into latent reasoning via Gumbel-Softmax, restoring LLMs’ exploratory capacity and enhancing their compatibility with Reinforcement Learning (RL). Building on this, we propose **Latent Reasoning Policy Optimization (LEPO)**, a novel framework that applies RL directly to continuous latent representations. Specifically, in rollout stage, LEPO maintains stochasticity to enable diverse trajectory sampling, while in optimization stage, LEPO constructs a unified gradient estimation for both latent representations and discrete tokens. Extensive experiments show that LEPO significantly outperforms existing RL methods for discrete and latent reasoning.¹

1 Introduction

Large Language Models (LLMs) have demonstrated remarkable reasoning capabilities via Chain-of-Thought (CoT) (Wei et al., 2022; Comanici et al., 2025; Guo et al., 2025). However, CoT is restricted to selecting one discrete token at each reasoning step, rather than retaining the distributional information inherent in the model’s internal states, which potentially induces information loss. To address this, recent research has shifted towards latent reasoning (Hao et al., 2024; Zhang et al., 2025). By utilizing continuous representations—such as hidden states (Hao et al., 2024) or weighted vocabulary embeddings (Zhang et al., 2025)—this paradigm

preserves richer information for more comprehensive reasoning (Zhu et al., 2025a).

However, due to the lack of stochastic sampling, existing latent reasoning methods (Hao et al., 2024; Zhang et al., 2025) yield an invariant output for a specific input. This deterministic process restricts the model’s exploration capabilities, potentially limiting its application. To better understand this restriction, we conceptually divide the model’s exploration capabilities into two levels: (1) **trajectory-level**: facilitating the sampling of diverse trajectories to empower the exploration of correct reasoning paths; (2) **step-level**: performing implicit search over multiple options within each step, thereby preventing early mistakes by maintaining diverse latent possibilities. Crucially, while current latent methods provide the step-level search through continuous representations, they sacrifice the trajectory-level exploration found in discrete approaches due to the lack of stochastic sampling.

To restore trajectory-level exploration, it is necessary to introduce stochasticity into latent reasoning. Technically, we achieve this by adding random perturbations to the continuous representation via Gumbel-Softmax (Jang et al., 2016; Wu et al., 2025). Our motivation experiments in Section 2.2 validate that simply injecting Gumbel noise effectively enhances exploration, as evidenced by increased entropy, a higher pass@32, and a shift in the problem difficulty distribution that converts intractable tasks into solvable ones. Furthermore, this restored exploration capability inherently aligns with the exploration-exploitation nature of reinforcement learning (RL), thereby unlocking the potential for superior performance gains.

Leveraging the stochasticity provided by Gumbel noise, we propose **Latent Reasoning Policy Optimization (LEPO)**, a novel framework that applies RL directly to the latent representation. LEPO effectively formulates latent reasoning as a learnable policy over continuous representation,

*Corresponding Authors (donghd66@gmail.com)

† Equal Contribution.

¹Code is available at <https://github.com/YuyanZhou/lepo>.

allowing the model to explore and optimize latent trajectories that are essential for solving complex tasks. During the rollout stage, we introduce controllable stochasticity via Gumbel-Softmax to sample diverse trajectories, thereby unlocking the exploration potential of latent reasoning. In the optimization stage, we extend the optimization objective from discrete tokens to latent representations, enabling unified gradient estimation over both latent reasoning steps and subsequent discrete answers. Extensive experiments show that LEPO significantly outperforms existing RL methods for discrete and latent reasoning on challenging mathematical and general-purpose benchmarks.

We summarize our major contributions as follows: (1) We conduct motivation experiments to show the superior exploration capability of stochastic latent reasoning. Our analysis reveals that latent reasoning with stochastic noise restores trajectory-level exploration, leading to higher entropy and facilitating more productive reinforcement learning. (2) We propose LEPO, an RL training framework tailored for stochastic latent reasoning. By applying policy optimization directly to continuous latent representations, LEPO effectively optimizes the model’s internal reasoning process. (3) Comprehensive experiments demonstrate that LEPO outperforms baseline RL frameworks for discrete and latent reasoning on mathematical and general-purpose benchmarks.

2 Preliminary and Motivation

2.1 From Discrete to Latent Reasoning

Standard LLMs perform reasoning via discrete token selection. Given a query \mathbf{q} and prior outputs $\mathbf{o}_{1:t-1}$, the next token \mathbf{o}_t is sampled from the vocabulary distribution $\pi \in \Delta^{|V|-1}$,

$$\mathbf{o}_t \sim \pi_t = \text{LLM}(\mathbf{q}, \mathbf{o}_{<t}) \in \Delta^{|V|-1}, \quad (1)$$

where $\Delta^{|V|-1}$ is the probability simplex over vocabulary V . In contrast, Latent Reasoning operates on the continuous latent token, most commonly realized as the vocabulary distribution (Zhang et al., 2025). It defines a **Latent Token** \mathbf{z}_t as:

$$\mathbf{z}_t := \pi_t = \text{LLM}(\mathbf{q}, \mathbf{E}_{<t}) \in \Delta^{|V|-1}, \quad (2)$$

where $\mathbf{E}_{<t}$ represents the input embeddings of previous steps. To proceed autoregressively, the input embedding \mathbf{E}_t for the next step is calculated as the expectation of token embeddings over the latent token \mathbf{z}_t :

$$\mathbf{E}_t := \sum_{k=1}^{|V|} \mathbf{z}_{t,k} \mathbf{e}_k, \quad (3)$$

where \mathbf{e}_k denote the embedding vector of the k -th token in the vocabulary.

2.2 Stochasticity Really Matters

As shown in Eq.2, unlike discrete sampling, current latent reasoning methods fail to sample diverse trajectories. Wu et al. (2025) recently proposed using Gumbel-Softmax to restore exploratory capacity. Building upon this insight, we conduct preliminary experiments to investigate whether this stochasticity can serve as a foundation for RL. We conduct preliminary experiments to investigate this exploration potential. First, we observe that stochastic latent reasoning exhibits higher entropy and pass@32 compared to discrete and deterministic latent methods, indicating superior exploration capability. Moreover, this enhanced exploration redistributes problems from the low-accuracy to the moderate-accuracy range. This phenomenon provides more solvable yet challenging tasks for the model, thereby yielding substantial performance gains during reinforcement learning.

Stochastic latent reasoning enhances exploration (Fig. 1a). Results show that stochastic latent reasoning maintains higher entropy, which naturally fosters diverse reasoning pathways, improving the discovery of correct solutions evidenced by higher pass@32.

Exploration mitigates problem hardness (Fig. 1b). By categorizing problems into base accuracy bins, we observe that the problem difficulty distribution is reshaped. Stochastic latent reasoning reduces the volume of intractable low-accuracy problems (in the 0.0-0.2 range) and converts them into solvable moderate-accuracy ones (in the 0.2-0.5 range).

Difficulty distribution shift benefits RL optimization (Fig. 1c). This redistribution enriches moderate-accuracy problems, providing solvable yet challenging tasks. These instances offer dense reward signals, unlocking the potential for RL to achieve superior performance.

3 Methods

Leveraging the exploration capability of stochastic latent reasoning, we design a novel RL-based framework, LEPO, for latent reasoning policy optimization.




Figure 1: Preliminary experiments using Qwen2.5-7B on the MATH500 dataset. (a) Stochastic latent reasoning enhances exploration (indicated by higher pass@32 and entropy) with comparable initial accuracy (mean@32). (b) Problem distribution across base accuracy bins, showing that stochastic latent reasoning shifts density from low-accuracy intervals to moderate-accuracy ones. (c) Performance gains after RL training, demonstrating that optimization is most effective on these moderate-accuracy problems.

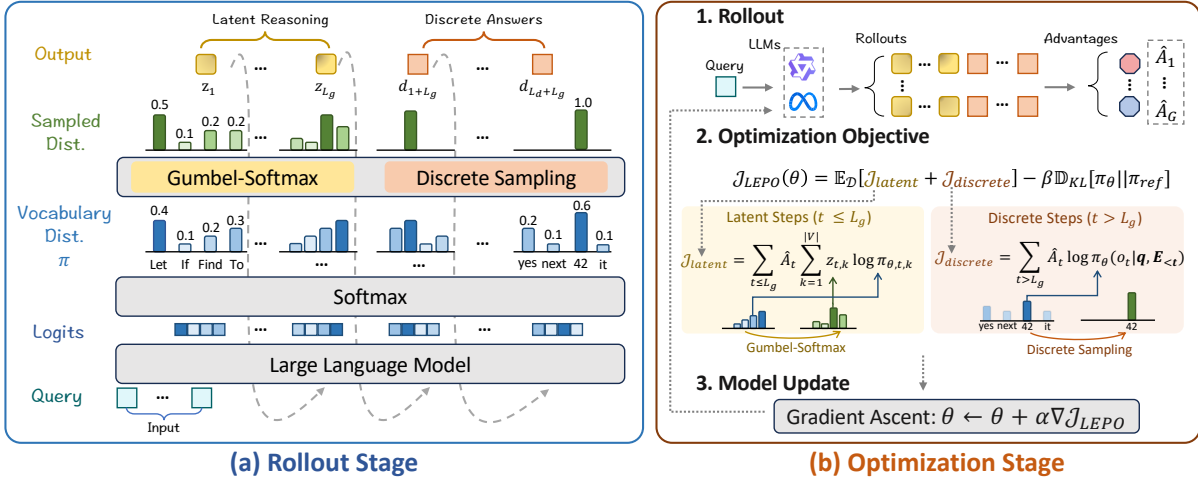


Figure 2: Illustration of the LEPO framework. (a) Rollout Stage: We employ Gumbel-Softmax for stochastic exploration of latent tokens, followed by discrete sampling for answers. (b) Optimization Stage: We optimize a unified objective that provides consistent supervision for both latent (soft targets) and discrete (hard targets) tokens based on rollout advantages.

As illustrated in Figure 2(a), we revitalize the rollout stage by incorporating a stochastic noise injection mechanism like Gumbel-Softmax, granting the LLM the flexibility to sample diverse trajectories. During the LEPO rollout stage, LLMs sequentially produce latent tokens with stochastic noise for reasoning and sample discrete tokens for readable answer generation. While in the optimization stage, beyond merely supervising discrete tokens, LEPO also integrates latent tokens into a unified gradient estimation framework. This unified optimization provides supervision to both latent and discrete tokens, thereby enhancing its effectiveness.

3.1 Sampling Diverse Trajectories via Gumbel-Softmax

Stochastic exploration is crucial in reinforcement learning. However, refer to Eq. 2 and 3, without a stochastic sampling mechanism, the latent reasoning steps become deterministic given the query \mathbf{q} , thus hindering exploration.

To enable trajectory-level exploration, following Wu et al. (2025), we inject stochasticity into the continuous representation. In our framework, this representation is embodied by the original distribution π_t over the vocabulary. Afterwards, we sample a stochastic latent token \mathbf{z}_t derived from π_t . The sampled \mathbf{z}_t is expected to meet three criteria: (1) **Validity**: It needs to be a valid distribution over the vocabulary; (2) **Latent Nature**: It is expected not to collapse into a one-hot distribution; (3) **Faithfulness**: It needs to preserve the core information of the original distribution.

Based on these criteria, we adopt the Gumbel-Softmax to obtain the derived distribution \mathbf{z}_t . Specifically, the k -th component of the stochastic latent token \mathbf{z}_t is calculated as:

$$\mathbf{z}_{t,k} = \frac{\exp((\log \pi_{t,k} + \epsilon_k)/\tau_g)}{\sum_{j=1}^{|\mathcal{V}|} \exp((\log \pi_{t,j} + \epsilon_j)/\tau_g)}, \quad (4)$$

where ϵ_k is the independent Gumbel noise derived from $\epsilon = -\log(-\log(u))$ where $u \sim \mathbf{U}(0, 1)$, τ_g

is the temperature to balance the latent nature and faithfulness. As $\tau_g \rightarrow 0$, \mathbf{z}_t approaches a one-hot distribution; conversely, as $\tau_g \rightarrow \infty$, \mathbf{z}_t approaches a uniform distribution. Distinct from other noise injection strategies, Gumbel-Softmax effectively approximates multinomial sampling, thereby preserving the distributional information more faithfully (Wu et al., 2025; Jang et al., 2016).

Leveraging this mechanism, we can sample diverse rollouts. Given a query q , we first autoregressively generate L_g stochastic latent tokens. At each step t , the latent token is sampled as:

$$\mathbf{z}_t = \text{GumbelSoftmax}(\text{LLM}(\mathbf{q}, \mathbf{E}_{<t}), \tau_g). \quad (5)$$

Crucially, this stochastic \mathbf{z}_t is then used to compute the input embedding \mathbf{E}_t via Eq. 3 for the next step. Following the latent reasoning phase, discrete tokens \mathbf{d} are sampled to generate the readable answer:

$$\mathbf{d}_t \sim \text{LLM}(\mathbf{q}, \mathbf{E}_{1:L_g}, \mathbf{d}_{<t}), \quad (6)$$

where $t = 1 + L_g, \dots, L_d + L_g$ and L_d is the length of answer. This process yields a hybrid trajectory consisting of explored latent reasoning steps followed by discrete answer tokens.

3.2 Unified Optimization Objective

Inspired by recent RL methods (Shao et al., 2024; Yu et al., 2025), we introduce a unified framework for effective joint optimization of latent and discrete tokens.

Given a query \mathbf{q} , we first sample a group of G rollouts. Based on a simple outcome-based reward r (1 for correct rollout and 0 otherwise), we calculate the advantages for the i -th rollout using normalized rewards:

$$\hat{A}_i = \frac{r_i - \text{mean}([r_1, \dots, r_G])}{\text{std}([r_1, \dots, r_G])}. \quad (7)$$

We then employ an on-policy RL algorithm to estimate the gradients. The optimization objective is divided into two parts based on the reasoning stage.

For Discrete Tokens ($t > L_g$): The objective follows the standard REINFORCE formulation. We maximize the log-likelihood of the selected token $o_{i,t}$ as:

$$\mathcal{J}_{\text{discrete}}^{(i)} = \sum_{t>L_g} \hat{A}_{i,t} \log \pi_{\theta}(o_{i,t} | \mathbf{q}, \mathbf{E}_{<t}). \quad (8)$$

For Latent Tokens ($t \leq L_g$): Since the latent token $\mathbf{z}_{i,t}$ is a continuous distributional vector rather than a discrete index, the standard probability mass function does not apply. To enable consistent supervision, we propose a soft surrogate objective. We treat the sampled stochastic vector $\mathbf{z}_{i,t}$ (from Eq. 4) as a "soft label" and maximize the alignment between the current policy π_{θ} and this label, weighted by the advantage:

$$\mathcal{J}_{\text{latent}}^{(i)} = \sum_{t \leq L_g} \hat{A}_{i,t} \sum_{k=1}^{|\mathcal{V}|} \mathbf{z}_{i,t,k} \log \pi_{\theta,k}(\mathbf{q}, \mathbf{E}_{<t}). \quad (9)$$

This term generalizes the discrete objective: if $\mathbf{z}_{i,t}$ were a one-hot vector, this equation would strictly recover the discrete policy gradient.

Unified Objective: Finally, let $\mathcal{J}_{\text{total}}^{(i)} = \mathcal{J}_{\text{latent}}^{(i)} + \mathcal{J}_{\text{discrete}}^{(i)}$ denote the total objective for the i -th trajectory. The overall LEPO optimization objective is computed as:

$$\mathcal{J}_{\text{LEPO}}(\theta) = \mathbb{E}_{\mathcal{D}} \left[\frac{1}{G} \sum_{i=1}^G \frac{1}{T_i} \mathcal{J}_{\text{total}}^{(i)} \right] - \beta \mathbb{D}_{\text{KL}}[\pi_{\theta} || \pi_{\text{ref}}], \quad (10)$$

where T_i is the length of the i -th trajectory, π_{ref} is the reference model for the KL-divergence regularization, and β is the coefficient. This unified framework ensures that both the latent reasoning process and the final answer generation are optimized towards high-reward outcomes.

4 Experiments

4.1 Experimental Settings

Baselines. We compare LEPO against a diverse array of baselines across three paradigms. In the realm of *inference-only methods*, we include CoT and CoT with latent inference; for *discrete-token RL*, we select GRPO (Shao et al., 2024) for comparison; while for *latent-token RL*, Soft Tokens (Butt et al., 2025) and HRPO (Yue et al., 2025) are employed as baselines. Additionally, we employ Gumbel-Softmax latent inference to the model trained by GRPO as a supplement.

Training and Evaluation Settings. We employ three prominent open-source models, i.e., Qwen2.5-7B, Qwen2.5-3B, and Llama-3.2-3B-Instruct as the base models to benchmark LEPO. The three models are trained on DAPO-MATH-17k (Yu et al., 2025) and then evaluated on GSM8k, MATH500, Minervamath, AIME2024, AIME2025,

Table 1: Main results on the reasoning benchmarks. We compare **LEPO** with discrete baselines (CoT, GRPO) and prior latent reasoning methods (Soft Tokens, HRPO) across three base models: Qwen2.5 (7B, 3B) and Llama3.2-3B-Instruct. We report Pass@1 and Pass@32 accuracy on six datasets. The results demonstrate that LEPO consistently outperforms other methods, achieving superior reasoning capabilities. The best and second-best results are highlighted in **bold** and underlined, respectively.

Method	GSM8K		MATH500		Minvervmath		AIME2024		AIME2025		AMC23		Average	
	Pass@1	Pass@32	Pass@1	Pass@32	Pass@1	Pass@32	Pass@1	Pass@32	Pass@1	Pass@32	Pass@1	Pass@32	Pass@1	Pass@32
Qwen2.5-7B														
CoT	44.43	<u>98.48</u>	35.15	93.20	14.10	54.04	4.48	23.33	3.44	30.00	25.62	85.00	21.20	64.01
CoT (w/ latent inference)	44.21	98.71	35.17	94.80	13.57	55.15	4.90	36.67	3.65	30.00	25.16	85.00	21.11	66.72
GRPO	89.47	97.57	74.51	93.40	33.98	57.72	10.10	<u>33.33</u>	5.10	33.33	47.89	90.00	43.51	<u>67.56</u>
GRPO (w/ latent inference)	<u>89.50</u>	98.03	<u>74.73</u>	93.40	<u>34.03</u>	<u>58.82</u>	9.38	26.67	6.15	33.33	49.14	92.50	<u>43.82</u>	<u>67.13</u>
Soft Tokens (Butt et al., 2025)	87.08	94.39	73.57	89.80	30.33	47.06	<u>10.94</u>	23.33	<u>7.29</u>	36.67	<u>50.47</u>	82.50	43.27	62.29
HRPO (Yue et al., 2025)	85.97	89.46	66.60	86.60	22.79	43.38	10.00	20.00	<u>7.29</u>	23.33	47.50	80.00	40.02	57.13
LEPO (Ours)	90.28	98.03	75.53	<u>93.80</u>	34.06	59.56	11.04	36.67	8.33	36.67	52.58	92.50	45.30	69.54
Qwen2.5-3B														
CoT	72.60	<u>98.18</u>	52.38	90.80	16.61	50.00	3.23	20.00	1.98	30.00	24.22	80.00	28.50	61.50
CoT (w/ latent inference)	72.93	98.48	52.40	<u>91.40</u>	16.91	<u>50.74</u>	3.33	23.33	1.98	<u>26.67</u>	25.11	82.50	28.78	62.19
GRPO	81.80	97.42	61.00	90.00	<u>22.84</u>	51.47	6.46	<u>26.67</u>	2.08	23.33	33.36	<u>87.50</u>	34.59	<u>62.73</u>
GRPO (w/ latent inference)	81.35	97.42	<u>62.59</u>	90.60	21.69	50.37	6.67	23.33	<u>3.10</u>	<u>26.67</u>	32.74	<u>87.50</u>	<u>34.69</u>	62.65
Soft Tokens (Butt et al., 2025)	77.56	91.74	55.26	82.60	17.13	33.82	4.48	20.00	1.15	23.33	27.42	77.50	30.50	54.83
HRPO (Yue et al., 2025)	77.63	89.46	56.40	84.00	16.54	33.09	<u>7.50</u>	23.33	3.33	20.00	<u>35.00</u>	77.50	32.73	54.56
LEPO (Ours)	82.94	97.57	63.23	91.60	24.22	<u>50.74</u>	7.81	30.00	<u>3.10</u>	<u>26.67</u>	35.08	90.00	36.06	64.43
Llama3.2-3B														
CoT	73.22	96.89	43.75	84.40	16.22	45.96	6.77	26.67	0.52	13.33	21.48	75.00	26.99	57.04
CoT (w/ latent inference)	73.32	97.42	43.59	<u>85.00</u>	15.80	45.96	7.19	30.00	<u>0.83</u>	16.67	21.64	77.50	27.06	<u>58.76</u>
GRPO	74.69	<u>97.27</u>	43.89	83.20	16.67	<u>47.06</u>	6.67	<u>33.33</u>	0.42	13.33	<u>24.14</u>	75.00	27.73	58.20
GRPO (w/ latent inference)	74.41	96.66	43.86	84.80	16.74	<u>47.06</u>	7.19	<u>33.33</u>	0.42	16.67	23.98	72.50	<u>27.77</u>	58.50
Soft Tokens (Butt et al., 2025)	75.27	93.03	<u>44.08</u>	81.60	<u>16.91</u>	43.38	4.79	30.00	0.42	10.00	20.55	72.50	27.00	55.10
HRPO (Yue et al., 2025)	<u>77.10</u>	90.30	39.60	79.60	15.81	41.18	10.00	26.67	0.42	6.67	17.50	67.50	26.73	51.97
LEPO (Ours)	77.29	96.97	46.51	86.20	17.65	48.16	<u>9.38</u>	40.00	0.96	16.67	27.03	77.50	29.80	60.92

Table 2: Out-of-distribution (OOD) evaluation on general reasoning tasks. We report Pass@1 and Pass@32 accuracy on GPQA-Diamond, ARC-C, and MMLU-STEM datasets.

Method	GPQA-Diamond		ARC-C		MMLU-STEM		Average	
	Pass@1	Pass@32	Pass@1	Pass@32	Pass@1	Pass@32	Pass@1	Pass@32
Qwen2.5-7B								
CoT	10.92	82.83	65.34	99.91	58.17	99.90	44.81	94.21
CoT (w/ latent inference)	10.65	83.33	64.51	99.91	57.72	<u>99.84</u>	44.29	94.36
GRPO	28.28	87.88	<u>87.08</u>	<u>98.98</u>	<u>78.75</u>	97.18	64.70	94.68
GRPO (w/ latent inference)	28.09	<u>88.89</u>	87.07	98.72	78.71	97.37	64.62	<u>94.99</u>
Soft Tokens (Butt et al., 2025)	<u>28.69</u>	81.31	86.67	95.48	76.63	93.24	63.99	90.01
HRPO (Yue et al., 2025)	27.84	78.79	86.12	94.71	76.40	92.77	63.45	88.76
LEPO (Ours)	29.96	90.91	87.20	98.89	79.05	97.40	65.40	95.73

and AMC23. To evaluate generalization and ensure capability retention, we also conduct OOD evaluation in three additional benchmarks: GPQA-Diamond (Rein et al., 2024), ARC-C (Clark et al., 2018), and MMLU-STEM (Hendrycks et al., 2021a). During the evaluation phase, we report Pass@1 and Pass@32 metrics to assess the quality of the learned policy and the coverage of the solution space, respectively. We utilize a maximum response length of 2,048, top-p=0.95, and top-k=30. The temperature is adjusted to 1.0 for training and 0.6 for evaluation.

4.2 Results

Main Results. We compare LEPO to all baselines and report the results in Table 1. Compared to the

discrete GRPO, LEPO improves Pass@1 by 1.76% and Pass@32 by 1.95% on average. Furthermore, compared to previous SOTA latent RL methods, Soft Tokens, it achieves gains of 3.46% and 7.56%. The training reward curves in Figure 5c show that LEPO exhibits superior convergence performance compared to GRPO. These results demonstrate that LEPO effectively enhances the model’s reasoning capabilities.

OOD Results. We conducted experiments on three out-of-distribution (OOD) tasks based on Qwen2.5-7B shown in Table 2. The results show that LEPO consistently achieves superior pass@1 performance across all OOD benchmarks, indicating its robust generalization capabilities after training.




Figure 3: Evolution of entropy. Comparison of entropy changes between LEPO and GRPO during the training process.




Figure 4: Comparison of token consumption during the training. The plot shows the change in token counts over training steps.

4.3 Training Dynamics and Analysis

Entropy. We track the evolution of entropy during the training process based on Qwen2.5-7B shown in Figure 3. It can be observed that the entropy exhibits a downward trend as training progresses. The entropy of LEPO remains consistently higher than that of GRPO, indicating that LEPO effectively maintains the exploratory nature of latent reasoning throughout training.

Token Efficiency. We report the evolution of token consumption for Qwen2.5-7B in Figure 4. Initially, the token counts for LEPO and GRPO are comparable. As the model trains, LEPO encourages the utilization of information-rich continuous representations, which results in a distinct reduction in token consumption compared to the baseline.

Latent Length L_g . We present the results for different latent lengths using Qwen2.5-7B in Fig. 5a. It demonstrates that LEPO achieves superior accuracy compared to the GRPO baseline across a broad range of latent lengths, especially with compact latent representations.

Gumbel Temperature τ_g . With the latent length fixed at 64, we examine how model performance varies with τ_g . The results in Fig. 5b indicate that LEPO consistently outperforms GRPO within the effective range ($\tau_g \in [0.3, 0.6]$), showing robust improvements under the Gumbel sampling configurations.

Noise Type. We modified the noise injection method to use Gaussian and Dirichlet noise. As shown in Fig. 5c, the methods with these types of noise have inferior convergence performance to that of GRPO and LEPO with Gumbel noise.

Case Study. We showcase a comparison of the reasoning steps and final answers produced by latent CoT, GRPO, and LEPO on an arithmetic task from GSM8K. As illustrated in Figure 6, without latent reasoning, the discrete method GRPO is misled by the one-hot tokens “initial cost” and overlooks “repair expense”, ultimately failing to reach the correct answer. While latent CoT leverages continuous latent tokens, without specialized training, it fails to deduct the ‘original price’ from the total, leading to an incorrect search path. In contrast, LEPO exhibits greedy-like behavior on most tokens by producing peaky distributions (e.g. token 4 “Calculate”), yet maintains meaningful uncertainty at key steps (e.g. token 8 “bought” and “repaired”). Consequently, LEPO achieves a superior trade-off between inference efficiency and path exploration, leading to more accurate and robust reasoning.

5 Related Work

Based on the training paradigms, latent reasoning methods can be categorized into three types: training-free, pre-training, and post-training. (1) **Training-free** latent methods use the weighted sum of vocabulary embeddings for latent reasoning without training (Zhang et al., 2025; Zhuang et al., 2025; Wu et al., 2025; Shi et al., 2025). (2) **Pre-training** latent methods attempt to incorporate continuous representations for direct reasoning during the pre-training phase, commonly via interleaving continuous and discrete tokens (Tack et al., 2025) or employing internal continuous CoT (Geiping et al., 2025; Zhu et al., 2025b). (3) **Post-training** latent methods can be categorized into SFT-based and RL-based approaches. SFT-based methods (Hao et al., 2024; Shen et al., 2025; Cheng and Van Durme, 2024) provide implicit supervision to latent representations by aligning discrete answer tokens with target ones. RL-based methods sample trajectories containing latent tokens and utilize rewards for supervision (Yue et al., 2025; Butt et al., 2025; Zheng and Lee, 2025; Tan et al., 2025). Butt et al. (2025) employs RLOO to optimize embeddings perturbed by Gaussian noise. Zheng and Lee (2025) utilizes the reparameterization trick for gradient estimation. Yue et al. (2025) mix latent and




Figure 5: Analysis of Hyperparameters and Training Dynamics using Qwen2.5-7B. (a) Performance comparison on MATH500 across varying latent lengths. (b) Performance comparison across varying Gumbel temperatures. (c) Training reward curves demonstrate that LEPO with Gumbel noise exhibits superior convergence compared to GRPO, as well as the Gaussian and Dirichlet noise variants.

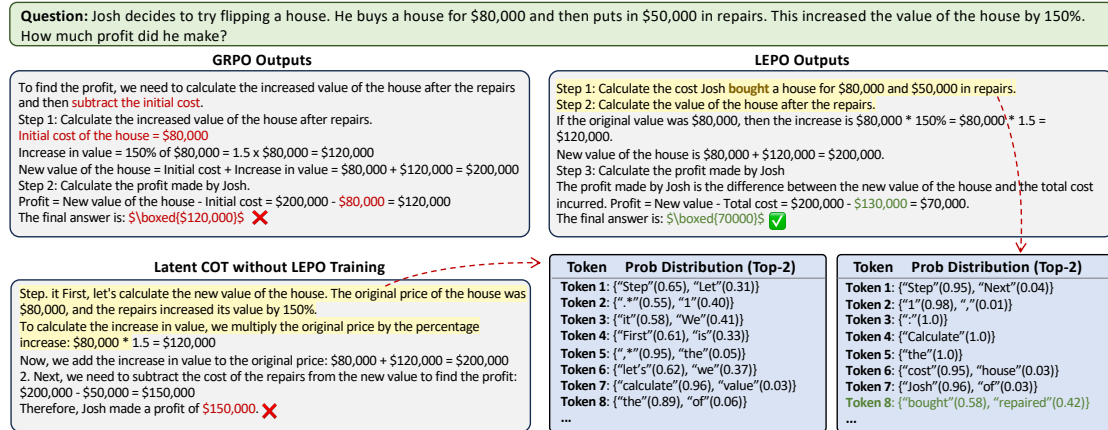


Figure 6: Qualitative comparison of GRPO, LEPO and latent COT on an arithmetic task from GSM8k. We highlight the latent tokens in yellow.

sampled discrete tokens at each step to introduce stochasticity. Our proposed **LEPO** falls into the category of latent RL methods. Distinct from these approaches, LEPO provides a unified optimization objective that supports the simultaneous optimization of both latent and discrete tokens within a single trajectory. Furthermore, instead of relying on Gaussian noise, we employ the Gumbel-Softmax mechanism to encourage the exploration of diverse trajectories. Empirically, our method achieves state-of-the-art performance.

6 Conclusion

In this work, we identify the limitation of deterministic inference in existing latent reasoning, which hinders trajectory-level exploration. By injecting controllable stochasticity via Gumbel-Softmax, we recover the model's capability to sample diverse reasoning paths. Building on this, we propose **LEPO**, a training framework that applies RL directly to continuous representations. Specifically, LEPO constructs a unified gradient estimation for both latent representations and discrete tokens, enabling the efficient optimization of internal reasoning policies.

7 Limitations

Despite the effectiveness of LEPO in enhancing reasoning capabilities, we acknowledge two primary limitations.

First, regarding interpretability, unlike discrete Chain-of-Thought methods that produce human-readable traces, our latent reasoning steps operate in a continuous space, which is not directly intelligible to humans. However, we argue that this is an inherent trade-off for leveraging the rich, dense information within vocabulary distributions. Furthermore, LEPO mitigates this issue by restricting the continuous representation strictly to the reasoning phase; the final answer generation phase utilizes standard discrete tokens, ensuring that the outcome remains fully readable. Additionally, for diagnostic purposes, the semantic meaning of the latent reasoning process can still be approximately interpreted by decoding the top-1 token from the latent distributions.

Second, concerning the generation paradigm, the current LEPO framework adopts a sequential "latent-first, discrete-later" structure. While this design effectively disentangles exploration from expression for mathematical benchmarks, it does not currently support the dynamic interleaving of

latent and discrete tokens (i.e., alternating between thinking and writing). We believe that extending LEPO to support such a hybrid, interleaved generation mode could further enhance flexibility for open-ended tasks, representing a promising direction for future work.

References

- Natasha Butt, Ariel Kwiatkowski, Ismail Labiad, Julia Kempe, and Yann Ollivier. 2025. Soft tokens, hard truths. [arXiv preprint arXiv:2509.19170](#).
- Jeffrey Cheng and Benjamin Van Durme. 2024. Compressed chain of thought: Efficient reasoning through dense representations. [arXiv preprint arXiv:2412.13171](#).
- Peter Clark, Isaac Cowhey, Oren Etzioni, Tushar Khot, Ashish Sabharwal, Carissa Schoenick, and Oyvind Tafjord. 2018. Think you have solved question answering? try arc, the ai2 reasoning challenge. [arXiv:1803.05457v1](#).
- Gheorghe Comanici, Eric Bieber, Mike Schaekermann, Ice Pasupat, Noveen Sachdeva, Inderjit Dhillon, Marcel Blistein, Ori Ram, Dan Zhang, Evan Rosen, and 1 others. 2025. Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities. [arXiv preprint arXiv:2507.06261](#).
- Jonas Geiping, Sean McLeish, Neel Jain, John Kirchenbauer, Siddharth Singh, Brian R Bartoldson, Bhavya Kailkhura, Abhinav Bhatele, and Tom Goldstein. 2025. Scaling up test-time compute with latent reasoning: A recurrent depth approach. [arXiv preprint arXiv:2502.05171](#).
- Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shiro Ma, Peiyi Wang, Xiao Bi, and 1 others. 2025. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. [arXiv preprint arXiv:2501.12948](#).
- Shibo Hao, Sainbayar Sukhbaatar, DiJia Su, Xian Li, Zhiting Hu, Jason Weston, and Yuandong Tian. 2024. Training large language models to reason in a continuous latent space. [arXiv preprint arXiv:2412.06769](#).
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2021a. Measuring massive multitask language understanding. [Proceedings of the International Conference on Learning Representations \(ICLR\)](#).
- Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. 2021b. Measuring mathematical problem solving with the math dataset. [arXiv preprint arXiv:2103.03874](#).
- Eric Jang, Shixiang Gu, and Ben Poole. 2016. Categorical reparameterization with gumbel-softmax. [arXiv preprint arXiv:1611.01144](#).
- David Rein, Betty Li Hou, Asa Cooper Stickland, Jackson Petty, Richard Yuanzhe Pang, Julien Dirani, Julian Michael, and Samuel R Bowman. 2024. Gpqa: A graduate-level google-proof q&a benchmark. In [First Conference on Language Modeling](#).
- Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, YK Li, Yang Wu, and 1 others. 2024. Deepseekmath: Pushing the limits of mathematical reasoning in open language models. [arXiv preprint arXiv:2402.03300](#).
- Zhenyi Shen, Hanqi Yan, Linhai Zhang, Zhanghao Hu, Yali Du, and Yulan He. 2025. Codi: Compressing chain-of-thought into continuous space via self-distillation. [arXiv preprint arXiv:2502.21074](#).
- Dachuan Shi, Abedelkadir Asi, Keying Li, Xiangchi Yuan, Leyan Pan, Wenke Lee, and Wen Xiao. 2025. Swireasoning: Switch-thinking in latent and explicit for pareto-superior reasoning llms. [arXiv preprint arXiv:2510.05069](#).
- Jihoon Tack, Jack Lanchantin, Jane Yu, Andrew Cohen, Ilia Kulikov, Janice Lan, Shibo Hao, Yuandong Tian, Jason Weston, and Xian Li. 2025. Llm pre-training with continuous concepts. [arXiv preprint arXiv:2502.08524](#).
- Wenhui Tan, Jiase Li, Jianzhong Ju, Zhenbo Luo, Jian Luan, and Ruihua Song. 2025. Think silently, think fast: Dynamic latent compression of llm reasoning chains. [arXiv preprint arXiv:2505.16552](#).
- Leandro von Werra, Younes Belkada, Lewis Tunstall, Edward Beeching, Tristan Thrush, Nathan Lambert, Shengyi Huang, Kashif Rasul, and Quentin Galouédec. 2020. Trl: Transformer reinforcement learning. <https://github.com/huggingface/trl>.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, and 1 others. 2022. Chain-of-thought prompting elicits reasoning in large language models. [Advances in neural information processing systems](#), 35:24824–24837.
- Chünhung Wu, Jinliang Lu, Zixuan Ren, Gangqiang Hu, Zhi Wu, Dai Dai, and Hua Wu. 2025. Llms are single-threaded reasoners: Demystifying the working mechanism of soft thinking. [arXiv preprint arXiv:2508.03440](#).
- Qiyang Yu, Zheng Zhang, Ruofei Zhu, Yufeng Yuan, Xiaochen Zuo, Yu Yue, Weinan Dai, Tiantian Fan, Gaohong Liu, Lingjun Liu, and 1 others. 2025. Dapo: An open-source llm reinforcement learning system at scale. [arXiv preprint arXiv:2503.14476](#).

Zhenrui Yue, Bowen Jin, Huimin Zeng, Honglei Zhuang, Zhen Qin, Jinsung Yoon, Lanyu Shang, Jiawei Han, and Dong Wang. 2025. Hybrid latent reasoning via reinforcement learning. [arXiv preprint arXiv:2505.18454](#).

Zhen Zhang, Xuehai He, Weixiang Yan, Ao Shen, Chenyang Zhao, Shuohang Wang, Yelong Shen, and Xin Eric Wang. 2025. Soft thinking: Unlocking the reasoning potential of llms in continuous concept space. [arXiv preprint arXiv:2505.15778](#).

Zhi Zheng and Wee Sun Lee. 2025. Soft-grpo: Surpassing discrete-token llm reinforcement learning via gumbel-reparameterized soft-thinking policy optimization. [arXiv preprint arXiv:2511.06411](#).

Hanlin Zhu, Shibo Hao, Zhiting Hu, Jiantao Jiao, Stuart Russell, and Yuandong Tian. 2025a. Reasoning by superposition: A theoretical perspective on chain of continuous thought. [arXiv preprint arXiv:2505.12514](#).

Rui-Jie Zhu, Zixuan Wang, Kai Hua, Tianyu Zhang, Ziniu Li, Haoran Que, Boyi Wei, Zixin Wen, Fan Yin, He Xing, and 1 others. 2025b. Scaling latent reasoning via looped language models. [arXiv preprint arXiv:2510.25741](#).

Yufan Zhuang, Liyuan Liu, Chandan Singh, Jingbo Shang, and Jianfeng Gao. 2025. Text generation beyond discrete token sampling. [arXiv preprint arXiv:2505.14827](#).

A More Experimental Setup

A.1 Prompt Examples

We provide the specific prompt templates used in our experiments. As shown in Figure 7, for mathematical reasoning tasks, the system instruction explicitly requires the model to reason step-by-step and enclose the final answer within `\boxed{}`. As shown in Figure 8, for multiple-choice benchmarks (i.e., GPQA-Diamond, ARC-C, and MMLU-STEM), the prompt is adapted to instruct the model to provide the specific option letter (A, B, C, or D) inside the box.

A.2 Implementation Details

For all experiments, we utilize the TRL framework (von Werra et al., 2020) as the training foundation. We employ the AdamW optimizer with a cosine learning rate scheduler, setting the warmup ratio to 0.03 and weight decay to 0.1. The optimizer hyperparameters are set to $\beta_1 = 0.9$ and $\beta_2 = 0.99$. Across all models, we set the learning rate to 1×10^{-6} and train for 1 epoch. The coefficient for the KL divergence penalty, β , is set to 1×10^{-3} . Regarding data sampling, we use a group

size of 8 for each query and a global batch size of 32. For all datasets, the maximum sequence length is set to 2048. We apply top- p sampling with $p = 0.95$ and top- k sampling with $k = 30$. The sampling temperatures are set to 1.0 during training and 0.6 for testing. All experiments are conducted with 8 Nvidia H20 GPUs. Training takes approximately 28 hours for Qwen2.5-7B, and 20 hours for Qwen2.5-3B and Llama-3.2-3B-Instruct.

For LEPO, we conduct a grid search over $L_g \in \{32, 64, 128, 256, 2048\}$ and $\tau_g \in \{0.3, 0.4, 0.5, 0.6, 0.7\}$, utilizing the MATH500 dataset for validation. Specifically, $L_g = 2048$ indicates that the entire response consists of latent tokens; in this scenario, to extract the final answer, we decode the top-1 token from the new distribution as the readable output. The optimal hyperparameters are identified as follows: for Qwen2.5-7B, $L_g = 64$ and $\tau_g = 0.5$; for Qwen2.5-3B, $L_g = 64$ and $\tau_g = 0.3$; and for the Llama3.2-3B-Instruct model, $L_g = 32$ and $\tau_g = 0.5$.

B Additional Results

B.1 Comparison to Deterministic Latent Reasoning with RL

In addition, we evaluated stochastic versus deterministic latent reasoning with RL. Since deterministic approaches lack randomness in the reasoning process, they cannot produce the valid supervision signals required for direct guidance, limiting optimization to the final answer. As evidenced by Table 3, LEPO achieves superior performance in both pass@1 and pass@32 compared to deterministic latent reasoning. This underscores the necessity of obtaining stochastic trajectories via noise injection, which serves the dual purpose of enabling RL-based supervision for latent reasoning and enhancing the model’s exploration capabilities.

B.2 Comparison to SFT-based Latent Reasoning Methods

In addition to comparing against strong RL baselines such as GRPO and HRPO, we also evaluated SFT-based latent reasoning methods, including Coconut (Hao et al., 2024) and CODI (Shen et al., 2025). Since the DAPO-MATH-17k dataset used in our main experiments lacks Chain-of-Thought (CoT) traces, we standardized the training on the MATH (Hendrycks et al., 2021b) dataset and evaluated performance on MATH500 and GSM8K. For all comparisons, we utilized Llama-3.2-3B-Instruct

Example Prompt For Mathematical Reasoning

```
<|im_start|>system
Please reason step by step, and put your final answer within \\boxed{{}}. <|im_end|>
<|im_start|>user
A robe takes 2 bolts of blue fiber and half that much white fiber. How many bolts in total does it take?
Please reason step by step, and put your final answer within \\boxed{{}}.<|im_end|>
<|im_start|>assistant
```

Figure 7: Example Prompt for Mathematical Reasoning in LEPO.

Example Prompt For GPQA-Diamond, ARC-C and MMLU-STEM

```
<|im_start|>system
Please reason step by step, and put your final answer within \\boxed{} with (A, B, C or D).<|im_end|>
<|im_start|>user
Question: Petrified palm trees are found in sedimentary rock near glaciers. The presence of the petrified palm
trees most likely provides evidence for which statement?

Options:
A. There was once more water in the area.
B. The area was once grassland.
C. The climate in the area was once tropical.
D. There are active faults in the area.
Please reason step by step, and put your final answer within \\boxed{{}}. with (A, B, C or D).<|im_end|>
<|im_start|>assistant
```

Figure 8: Example Prompt for GPQA-Diamond, ARC-C and MMLU-STEM in LEPO.

as the backbone model. As shown in Table 4, LEPO significantly outperforms both Coconut and CODI. This demonstrates the effectiveness of applying reinforcement learning to supervise latent reasoning processes, exhibiting its consistent advantages over prior SFT-based latent methods.

B.3 Analysis on More Base Models

In addition to the analysis of Qwen2.5-7B presented in Figures 3, 4 and 5, we further investigate the training dynamics, entropy, and token efficiency of Qwen2.5-3B and Llama-3.2-3B-Instruct.

Training Dynamics. As illustrated in Figures 9a and 10a, we compare the training dynamics of LEPO and GRPO on the Qwen2.5-3B and Llama-3.2-3B-Instruct models. The results indicate that while both LEPO and GRPO effectively improve model performance, LEPO achieves superior final convergence.

Entropy. As shown in Figure 9b, we observe that the training entropy for Qwen2.5-3B gradually decreases, with LEPO consistently maintaining higher entropy than GRPO. Conversely, as depicted in Figure 10b, the entropy for Llama-3.2-3B-Instruct initially decreases before rising; notably, LEPO’s entropy remains consistently higher than that of GRPO throughout. This demonstrates that stochastic latent reasoning possesses greater ex-

ploratory capability compared to discrete representations and maintains this heightened exploration during the training process.

Token Efficiency. As illustrated in Figure 9c, while the reasoning length of Qwen2.5-3B gradually increases during training, LEPO utilizes continuous representations to reduce the number of tokens compared to GRPO. Similarly, in Figure 10c, although the token count for Llama-3.2-3B-Instruct gradually decreases, LEPO consistently utilizes fewer tokens. These results highlight the superiority of LEPO over GRPO in terms of token efficiency.

C The Use of LLMs

During the drafting phase of this work, we made use of LLMs to assist with text composition and editing. These models served primarily to enhance the flow of the narrative, adjust phrasing, and ensure grammatical correctness. Crucially, the intellectual contributions—including all hypotheses, experimental setups, and derived insights—originate solely from the authors, who retain complete accountability for the manuscript’s integrity.

Table 3: Comparison of stochastic versus deterministic latent reasoning on Qwen2.5-7B. We evaluate the performance using Pass@1 and Pass@32 metrics. LEPO significantly outperforms the deterministic baseline, underscoring the necessity of noise injection for enabling effective RL supervision and enhancing exploration capabilities.

Method	GSM8K		MATH500		Minervamath		AIME2024		AIME2025		AMC23		Average	
	Pass@1	Pass@32	Pass@1	Pass@32	Pass@1	Pass@32	Pass@1	Pass@32	Pass@1	Pass@32	Pass@1	Pass@32	Pass@1	Pass@32
Qwen2.5-7B														
Deterministic Latent	88.75	94.62	73.84	91.80	32.49	51.10	8.04	26.67	6.54	23.33	47.27	85.00	42.82	62.09
LEPO (Ours)	90.28	98.03	75.53	93.80	34.06	59.56	11.04	36.67	8.33	36.67	52.58	92.50	45.30	69.54




Figure 9: Training analysis of LEPO versus GRPO on Qwen2.5-3B. (a) Evolution of training rewards, where LEPO outperforms GRPO. (b) Policy entropy trends; despite the initial drop and subsequent rise, LEPO maintains consistently higher entropy than GRPO. (c) Evolution of generated token counts, demonstrating the superior token efficiency of LEPO compared to the baseline.

Table 4: Comparison with SFT-based latent reasoning methods (Coconut and CODI) on GSM8K and MATH500. All models utilize Llama-3.2-3B-Instruct as the backbone and are trained on the MATH dataset.

Method	GSM8K	MATH500
Llama3.2-3B		
Coconut	33.51	-
CODI	69.98	37.80
LEPO (Ours)	76.57	46.80




Figure 10: Training analysis of LEPO versus GRPO on Llama-3.2-3B-Instruct. (a) Evolution of training rewards, where LEPO outperforms GRPO. (b) Policy entropy trends; despite the initial drop and subsequent rise, LEPO maintains consistently higher entropy than GRPO. (c) Evolution of generated token counts, demonstrating the superior token efficiency of LEPO compared to the baseline.