

Latent-SFT: LLM Latent Reasoning as Chain of Superposition

Jingcheng Deng^{1,2}, Liang Pang^{1*}, Zihao Wei^{1,2}, Shicheng Xu^{1,2}, Zenghao Duan^{1,2},
Kun Xu, Yang Song, Huawei Shen^{1,2}, Xueqi Cheng^{1,2}

¹State Key Laboratory of AI Safety,

Institute of Computing Technology, Chinese Academy of Sciences, Beijing, China

²University of Chinese Academy of Sciences, Beijing, China

Emails: dengjingcheng23s@ict.ac.cn, pangliang@ict.ac.cn

Abstract

Latent reasoning offers a computation-efficient alternative to Chain-of-Thought but often suffers from performance degradation due to distributional misalignment and ambiguous chain definitions. Ideally, latent reasoning should function as a superposition of multiple reasoning paths. To realize this, we introduce Latent-SFT, a unified framework addressing challenges at three levels: token, chain, and learning. First, we define the Latent-Vocab to constrain hidden states within the pre-trained vocab-space. Second, we construct the Latent-Chain via Induction-Supervision Masking to ensure semantic compactness and sufficiency. Third, we employ Latent-Optim with stochastic Gumbel-Softmax to guide the model toward generalizable solutions. Empirical results demonstrate that Latent-SFT consistently outperforms explicit SFT across six mathematical benchmarks (e.g., GSM8k, AIME24) while achieving a $2.7\times$ to $5.5\times$ reduction in reasoning length. Analysis confirms that our method effectively captures a superposition of diverse reasoning trajectories rather than merely compressing a single path.

1 Introduction

Large language models (LLMs) have shown strong reasoning abilities (Zhao and Zhang, 2024) across diverse tasks (Cao and Zhao, 2025; Wei et al., 2025; Deng et al., 2025a,b). A key driver of this success is chain-of-thought prompting (Wei et al., 2022; Xu et al., 2024; Yao et al., 2023), which solves complex problems by generating intermediate steps in natural language. However, current reasoning models (Team, 2025; DeepSeek-AI, 2025) rely on long chains of tokens (Xu et al., 2026; Wei et al., 2026), leading to redundancy and high computational cost, and motivating a growing body of work on efficient reasoning (Zhang et al., 2025a; Xia et al., 2025; Arora and Zanette, 2025). Latent Reasoning is a




Figure 1: Comparison of Latent and Explicit Reasoning. Latent reasoning operates as a superposition state, where each step carries a higher entropy. Therefore, it can complete the reasoning process and derive the correct answer using fewer tokens.

promising paradigm that significantly reduces the length of reasoning chains by shifting computation to a continuous latent space. This paradigm theoretically enhances information bandwidth and exploration efficiency (Zhu et al., 2025), yet current methods frequently suffer from severe performance degradation compared to explicit reasoning (Hao et al., 2024).

Ideally, latent reasoning should function as a superposition of multiple reasoning paths, enabling the model to explore diverse logical trajectories simultaneously within a latent space. However, realizing this potential is impeded by three hierarchical flaws in current methodologies. First, at the **latent token level**, existing approaches rely on raw, unconstrained hidden states (Hao et al., 2024; Cheng and Durme, 2024). These high-dimensional vectors suffer from severe distributional shifts relative to the model’s vocab-space, making them difficult for the model to interpret. Second, at the **latent chain level**, current definitions are ambiguous; they typically assume the latent chain need only decode the final answer (Shen et al., 2025). This "black box" assumption neglects intermediate semantic fidelity, failing to capture the step-by-step logic required for

* Corresponding author.

complex deduction. Third, at the **latent learning level**, optimization is inefficient. Models are updated solely via indirect backpropagation from the final explicit supervision (Hao et al., 2024), lacking the step-by-step guidance necessary to learn stable reasoning trajectories.

To systematically address these challenges, we introduce **Latent-SFT**, a comprehensive framework that unifies token-level representation, chain-level definition, and learning strategies. For token-level representation, we formally define the **Latent-Vocab** as a superposition within the vocab-space (Section 4). This formulation aligns latent tokens with the model’s inherent vocabulary distribution. For the chain-level definition, we introduce the **Latent-Chain** (Section 5), constructed via Induction-Supervision Masking to satisfy semantic compactness, compatibility, and sufficiency. Finally, regarding the learning strategy, to enable the LLM to autonomously generate these latent chains, we propose **Latent-Optim**. Within this framework, we address the static nature of the constructed latent chains by employing stochastic Gumbel-Softmax optimization. This technique injects Gumbel noise to approximate sampling, acting as an implicit Hessian regularizer (see Theorem 1 in Appendix G) that guides the model toward flatter, more generalizable minima.

As illustrated in Figure 1, our model realizes latent reasoning as a chain of superposition: each latent step is a high-entropy distribution that retains weighted support for several candidate trajectories at once, and this entropy decreases as the reasoning converges on a discrete answer. Empirical results demonstrate that Latent-SFT consistently outperforms explicit SFT models across six diverse mathematical reasoning benchmarks. Notably, it achieves a $2.7\times$ to $5.5\times$ reduction in reasoning chain length while surpassing prior latent reasoning approaches. Further analysis reveals that its latent reasoning is not merely a compression of a single chain, but rather a superposition of multiple reasoning chains.

2 Related Work

Current research remains entangled in debates over the reasonable representation of latent tokens, while the critical dimensions of latent chain definition and optimization strategies have been largely overlooked. Consequently, the definition of the latent token serves as the primary differentiator in this field. We categorize existing research based on

this pivotal factor.

Unconstrained Representations. Pioneering approaches, such as COCONUT (Hao et al., 2024) and CCOT (Cheng and Durme, 2024), adopted a direct approach by utilizing the last-layer hidden states as latent tokens. While computationally straightforward, these methods rely on raw, high-dimensional vectors that are unconstrained by the model’s vocabulary distribution. Although subsequent works like CODI (Shen et al., 2025) and PCCOT (Wu et al., 2025a) attempted to mitigate this misalignment via MLP mapping or distillation, they fundamentally remain bound to the continuous hidden state paradigm. Consequently, without strict constraints aligned with the pre-trained semantic manifold, these methods often face optimization instability and limited performance gains compared to explicit reasoning.

Constrained Representations. To address the distributional gap, recent studies have introduced structural constraints. CoLaR (Tan et al., 2025) represents a step forward by modeling latent tokens via Gaussian distributions, achieving improved alignment through a dedicated latent head. Similarly, methods like Soft Thinking (Zhang et al., 2025b), Mixture-of-Input (MoI) (Zhuang et al., 2025), and recent vocabulary-space superposition approaches (Zeng et al., 2025; Butt et al., 2025) attempt to bridge the gap by constructing latent tokens as linear combinations of vocabulary embeddings. However, their performance gains on reasoning tasks remain marginal. In contrast, we formally define the Latent-Vocab and Latent-Chain to enable simultaneous multi-path exploration, thereby significantly shortening reasoning chains while surpassing existing baselines.

3 Latent-SFT

Latent-SFT consists of three coupled components. First, **Latent-Vocab** defines each latent token as a simplex-constrained vocabulary distribution α and its corresponding superposition embedding $\mathbf{z} = \mathbf{E}\alpha$. Second, **Latent-Chain** (Stage 1, top panel of Figure 2) converts explicit CoT trajectories into static latent-chain targets $\{(\alpha_i, \mathbf{z}_i)\}_{i=1}^N$ through induction-supervision masking. Third, **Latent-Optim** (Stage 2) trains the LLM to generate these latent-chain distributions directly. At inference time, the construction encoder is not used; the trained LLM autoregressively produces latent-




Figure 2: Overview of Latent-SFT. **Latent-Vocab** defines each latent token as a vocabulary distribution α_i and its superposition embedding z_i . **Stage 1** (top panel, **Latent-Chain**) constructs static latent-chain targets $\{(\alpha_i, z_i)\}$ from explicit CoT trajectories through information compression, representation alignment, and supervised decoding under Induction–Supervision Masking. **Stage 2** (bottom panel, **Latent-Optim**) trains the LLM to generate such latent chains directly using KL and CE supervision. At inference time, only the trained LLM is used to produce latent-token distributions and decode the final answer.

token distributions and then decodes the final answer. The complete algorithmic flow is summarized in Algorithm 1.

4 Latent-Vocab: Superposition of Vocab-Space

Prior studies have explored both unconstrained hidden states and several constrained alternatives for latent tokens, yet it remains unclear what kind of constraint best aligns latent reasoning with the pre-trained language-model interface. We therefore conduct two preliminary analyses addressing the fundamental question: *what is the reasonable representation for latent tokens?*

4.1 Distributional Gap in Vector Spaces

Most existing studies (Shen et al., 2025; Wu et al., 2025a) follow the COCONUT design and use last-layer hidden states directly as latent tokens. However, statistical analysis in representative LLMs reveals a critical discrepancy between these vector spaces.

Observation 1. *Significant distributional divergence exists between the last-layer hidden states and the token embedding space.* As illustrated in




Figure 3: Visualization of last-layer hidden states and token embeddings of the LLaMA-3.2-1B-Instruct model on the GSM8k dataset. Results are shown from both global and token-specific perspectives to highlight the substantial distributional differences. In addition, we report the statistical differences between the two distributions using the Fréchet Inception Distance (FID) and the squared Maximum Mean Discrepancy (MMD²). We also estimate cosine similarity via random sampling.

Figure 3 and Figure 7, PCA projections demonstrate a substantial separation between the manifolds of hidden states and token embeddings. Quantitative analysis confirms this divergence, showing significant disparities in first- and second-order statistics (mean and variance), as well as large inter-distribution distances. From a model optimization perspective, shallow parameters (e.g., RMSNorm) are calibrated exclusively on the token embedding distribution during pre-training. Consequently, us-

ing raw hidden states as latent tokens can introduce distribution shift and training instability. We further find that this gap is not reduced, but enlarged, in the released COCONUT checkpoint (Appendix C), suggesting that hidden-state latent tokens may remain misaligned with the pretrained vocabulary interface after latent-reasoning training.

4.2 Structured Anisotropy of Vocab-Space

Building on Observation 1, we posit that latent tokens should preserve the pretrained geometry of vocab-space. Since LLMs are trained through discrete next-token prediction, their semantic interface is not an unconstrained continuous Gaussian space, but a vocabulary distribution mapped through token embeddings and the language modeling head. To examine this geometry, we analyze the singular value spectra of token embedding matrices across various LLMs.

Observation 2. *The vocabulary embedding space is structured and anisotropic rather than an unconstrained isotropic vector space.* The singular value decay curves in Figure 8 (Appendix D) show that spectral energy is distributed unevenly across directions. Rather than implying a low-dimensional collapse, this suggests that latent tokens should inherit the structured vocabulary interface induced by discrete language modeling, instead of being modeled as free hidden states or Gaussian latent variables.

These findings imply that latent tokens should be represented through the same vocabulary interface used during pretraining. Motivated by this insight, we propose the following definition.

Definition 1 (*Superposition Embedding*). We define a latent token $z \in \mathbb{R}^d$ as a simplex-constrained superposition over vocab-space. Formally, let $\mathbf{E} = [e_1, e_2, \dots, e_V] \in \mathbb{R}^{d \times V}$ denote the token embedding matrix. A latent token z is expressed as a probability-weighted mixture of token embeddings:

$$z = \mathbf{E}\alpha = \sum_{i=1}^V \alpha_i e_i, \quad (1)$$

where $\alpha = [\alpha_1, \dots, \alpha_V]^\top \in \Delta^{V-1}$ is a token-probability vector. We do not require a unique linear decomposition; the key is a vocabulary-aligned parameterization compatible with the LM head and the pretrained discrete-token interface.

5 Latent-Chain: Forming Latent Chain via Induction-Supervision Masking

Having established Latent-Vocab, we address the next core question: *what are the requisite properties of the higher-level Latent-Chain?* To avoid treating latent reasoning as a black box optimized only for the final answer, we require three semantic criteria: (1) Semantic Compactness, compressing multiple explicit tokens into one latent representation; (2) Semantic Compatibility, adhering to the vocab-space defined in Section 4; and (3) Semantic Sufficiency, retaining enough information for subsequent reasoning. Figure 2 implements these criteria through information compression, representation alignment, and supervised decoding.

Information Compression Process (Semantic Compactness). We insert special tokens into the explicit reasoning sequence to segment and capture intermediate reasoning information. Given an input $X = [Q, \langle \text{think} \rangle, C, \langle / \text{think} \rangle]$, we define a segmentation function $\text{Seg}(\cdot)$ that partitions the explicit reasoning chain C into N subsegments, denoted as $C = \{S_i\}_{i=1}^N$. This segmentation can be governed either by a fixed compression ratio r or by semantic boundaries; in our experiments, we use the fixed-length version detailed in Appendix H.1. Following each subsegment S_i , we insert a special latent placeholder L_i , which is a newly added special token in the vocabulary. Crucially, we apply a *Latent Token Induction Mask* (LTIM) to restrict the attention of L_i exclusively to the preceding explicit reasoning context (see Figure 2). This design ensures that the semantic content encoded by each latent token represents a cumulative summary of the reasoning process up to that point.

Representation Alignment Process (Semantic Compatibility). For each inserted placeholder L_i , we extract its hidden state $\mathbf{h}_i \in \mathbb{R}^d$ from the model under LTIM masking. To strictly enforce **Semantic Compatibility**, we map $\mathbf{h}_i$ to a vocabulary distribution through the language modeling head and reconstruct it as a mixture of token embeddings. Let $\mathbf{E} = [e_1, \dots, e_V] \in \mathbb{R}^{d \times V}$ denote the token embedding matrix and $\mathbf{W} \in \mathbb{R}^{d \times V}$ be the language modeling head (where $\mathbf{W} = \mathbf{E}$ under weight tying). We compute the vocabulary logits and the resulting probability distribution as:

$$\ell_i = \mathbf{W}^\top \mathbf{h}_i, \quad \alpha_i = \text{softmax}\left(\frac{\ell_i}{\tau_{\text{enc}}}\right) \in \Delta^{V-1}, \quad (2)$$

where $\tau_{\text{enc}} > 0$ is the encoder temperature; we set $\tau_{\text{enc}} = 1$ in all experiments. The final latent representation is then obtained via linear interpolation:

$$\mathbf{z}_i = \mathbf{E}\boldsymbol{\alpha}_i = \sum_{v=1}^V \alpha_{i,v} \mathbf{e}_v. \quad (3)$$

This operation instantiates Definition 1, guaranteeing that $\mathbf{z}_i$ resides within the convex hull of the embedding space $\{\mathbf{e}_v\}_{v=1}^V$. Consequently, latent tokens remain confined to the intrinsic geometry of the LLM’s vocabulary manifold. During the subsequent optimization phase, the distribution $\boldsymbol{\alpha}_i$ serves as the supervision target (via KL divergence) for the latent positions.

Supervised Decoding Process (Semantic Sufficiency). Given the latent representations $\{\mathbf{z}_i\}_{i=1}^N$ from the alignment step, we use each $\mathbf{z}_i$ to decode the content from the $(i + 1)$ -th explicit reasoning subsegment through to the final answer. Let the mixed prefix up to step i be $\Pi_i = [Q, \langle \text{think} \rangle, \mathbf{z}_1, S_2, \mathbf{z}_2, \dots, S_i, \mathbf{z}_i]$. Here S_1 is replaced by $\mathbf{z}_1$, while $S_2, \dots, S_i$ remain visible. Under the Latent Token Supervision Mask (LTSuM), tokens in $\mathcal{Y}_i = [S_{i+1}, \dots, S_N, \langle \text{think} \rangle, \text{Answer}]$ are decoded conditioned on Π_i and the causal prefix within $\mathcal{Y}_i$, while attention to future latent tokens and non-causal target tokens is blocked. Denoting by $\mathcal{J}_i$ the index set of tokens in $\mathcal{Y}_i$, the supervised decoding objective is

$$\mathcal{L}_{\text{sup}} = \frac{1}{N} \sum_{i=1}^N \frac{1}{|\mathcal{J}_i|} \sum_{t \in \mathcal{J}_i} \left(-\log p_{\theta}(x_t \mid \Pi_i; \text{LTSuM}) \right). \quad (4)$$

Having defined the three semantic properties above, we next describe how to optimize the latent-token encoder and the target LLM to construct high-quality latent chains in practice.

Latent-Chain Construction Training. The strict LTIM/LTSuM constraints can destabilize naive end-to-end SFT. We therefore adopt a progressive schedule: we first optimize the latent-token encoder, then the LLM backbone, and finally fine-tune both components jointly under $\mathcal{L}_{\text{sup}}$. This EM-style staging lets each module adapt to latent-token induction and supervised decoding before full co-adaptation, stabilizing training under the distributional shift introduced by latent tokens. Full

details are provided in Appendix H, and a breakdown of the resulting training cost in Appendix E.

6 Latent-Optim: Training LLMs for Latent Reasoning

Having constructed high-quality Latent-Chains in Section 5, we next train the LLM to generate them autonomously during reasoning. Latent-Optim provides explicit supervision for intermediate latent positions, avoiding reliance on indirect backpropagation from the final answer alone while improving generalization.

We concatenate latent chains with the explicit context to form the training sequence $X = [Q, \langle \text{think} \rangle, \mathbf{z}_1, \dots, \mathbf{z}_N, \langle \text{think} \rangle, \text{Answer}]$. The explicit segments S_i are removed in this stage because they are used only to supervise the latent-token encoder during Stage 1; here each $\mathbf{z}_i = \mathbf{E}\boldsymbol{\alpha}_i$ is a base latent target constructed by the trained encoder. A straightforward optimization approach is to minimize the KL divergence between the static latent chain $\boldsymbol{\alpha}_t$ and the model’s predicted vocabulary distribution $\mathbf{q}_t = \text{softmax}(\mathbf{W}^\top \mathbf{h}_t)$ at latent position t , combined with the standard cross-entropy loss for explicit answer tokens S_{exp} :

$$\mathcal{L}_{\text{static}} = \frac{1}{|S_{\text{lat}}|} \sum_{t \in S_{\text{lat}}} \text{KL}(\boldsymbol{\alpha}_t \parallel \mathbf{q}_t) + \frac{1}{|S_{\text{exp}}|} \sum_{t \in S_{\text{exp}}} (-\log q_t[y_t]). \quad (5)$$

However, directly supervising the LLM with static latent chains $\boldsymbol{\alpha}_t$ can lead to overfitting on the specific artifacts of the construction encoder, thereby limiting the model’s generalization capability.

To address this, we propose a **Stochastic Latent Optimization** strategy. Instead of treating the latent targets as immutable labels, we view them as parameters of a generative distribution. During each training epoch, we dynamically sample a perturbed latent distribution using the Gumbel-Softmax trick (Jang et al., 2017; Wu et al., 2025b). Specifically, for a latent position $t \in S_{\text{lat}}$, let $\boldsymbol{\alpha}_t$ be the static latent chain. We construct a stochastic distribution $\tilde{\mathbf{p}}_t$ and its corresponding latent input $\tilde{\mathbf{z}}_t$ by injecting Gumbel noise:

$$\begin{aligned} \tilde{\mathbf{p}}_t &= \text{softmax} \left(\frac{\log \boldsymbol{\alpha}_t + \lambda \mathbf{g}_t}{\tau_g} \right), \\ \tilde{\mathbf{z}}_t &= \mathbf{E}\tilde{\mathbf{p}}_t, \\ \mathbf{g}_t &\sim \text{Gumbel}(0, 1), \end{aligned} \quad (6)$$

where λ controls the noise strength and τ_g is the Gumbel-Softmax temperature. We set $\lambda = 1$ and $\tau_g = 1$ in all reported experiments. The noise vector $\mathbf{g}_t$ is resampled independently at every training step and applied to both the KL supervision target $\tilde{\mathbf{p}}_t$ and the teacher-forced latent input $\tilde{\mathbf{z}}_t$.

Consequently, the final optimization objective is reformulated as the expectation over this noise:

$$\mathcal{L}_{\text{auto}}(\theta_{\text{llm}}) = \mathbb{E}_{\mathbf{g}_{1:N}} \left[\frac{1}{|S_{\text{lat}}|} \sum_{t \in S_{\text{lat}}} \text{KL}(\tilde{\mathbf{p}}_t \parallel \mathbf{q}_t) + \frac{1}{|S_{\text{exp}}|} \sum_{t \in S_{\text{exp}}} (-\log q_t[y_t]) \right], \quad (7)$$

where $\mathbf{q}_t$ is evaluated under the teacher-forced stochastic latent prefix (and therefore depends on the preceding sampled latent inputs). By minimizing the expectation over the Gumbel noise, we effectively smooth the optimization landscape. As we formally prove in **Theorem 1** (see Appendix G), this stochastic injection acts as an implicit regularizer that penalizes the sharpness of the loss Hessian, thereby guiding the model toward flatter minima and enhancing generalization across diverse reasoning paths.

Inference. At inference time, the Latent-Chain construction encoder is discarded. Given a question Q , the trained LLM first enters the latent reasoning stage after `<think>`. At each step t , the model produces a hidden state $\mathbf{h}_t$ and maps it to a vocabulary distribution $\alpha_t = \text{softmax}(\mathbf{W}^\top \mathbf{h}_t)$. We first decode a candidate discrete token from α_t ; if it is `</think>`, we append the discrete stop token and terminate latent mode without appending a latent token at this position. Otherwise, we feed the corresponding superposition embedding $\mathbf{z}_t = \mathbf{E}\alpha_t$ as the next latent input. Thus, the hidden state is used only to parameterize a vocabulary distribution through the LM head; we do not feed unconstrained hidden states as latent tokens. The number of latent steps T is determined dynamically by the first generated `</think>` token. If no stop token is generated by T_{max} , the sample is counted as unsuccessful. We report $\#L$ as the number of generated latent superposition tokens.

7 Experiments

7.1 Experimental Setups

Datasets. To comprehensively assess the effectiveness of our method, we conducted experiments on both low-difficulty mathematical reasoning datasets (GSM8k-Aug (Deng et al., 2024), GSM-Hard (Gao et al., 2023), SVAMP (Patel et al., 2021), and MultiArith (Roy and Roth, 2015)) and high-difficulty datasets (MATH-500 (Hendrycks et al., 2021) and AIME24). For the low-difficulty tasks, we trained on the GSM8k-Aug training split to ensure fair comparison with prior work. For the high-difficulty tasks, due to limited computational resources, we used a subset of Open-R1 (Hugging Face, 2025) containing explicit reasoning chains shorter than 4k tokens, resulting in 57,544 training samples. This restriction inevitably constrained the performance of our method.

Baselines and Models. The baselines considered in our experiments include COT-SFT (Wei et al., 2022), iCOT (Deng et al., 2024), COCONUT (Hao et al., 2024), CODI (reproduced) (Shen et al., 2025), and CoLaR (Tan et al., 2025). To align model capabilities with task complexity, we strictly match base models to dataset difficulty. For low-difficulty tasks, we adopt LLaMA-3.2-1B-Instruct and LLaMA-3.2-3B-Instruct (Dubey et al., 2024), while for high-difficulty tasks, we employ DeepSeek-R1-Distill-Qwen-7B (DeepSeek-AI, 2025). We report Pass@1, reasoning length ($\#L$), and their ratio Pass@1/ $\#L$; full evaluation protocol and implementation details are provided in Appendix I. Throughout, Latent-SFT(r) denotes a model trained with compression ratio r : every r consecutive explicit CoT tokens are compressed into a single latent token (Appendix H.1). Latent-SFT(2) and Latent-SFT(4) therefore compress the explicit chain by $2\times$ and $4\times$, and CoLaR- k follows the analogous convention.

7.2 Main Results

Performance on Low-Difficulty Tasks. Table 1 compares Latent-SFT with baselines under different compression ratios r . Latent-SFT consistently improves Pass@1 and shortens reasoning chains relative to latent-reasoning baselines, and generalizes well to all three out-of-distribution datasets. On the 1B model, Latent-SFT(2) improves average Pass@1 by 2.0 percentage points while reducing chain length by approximately $5.4\times$, outperform-

	GSM8k-Aug		GSM-Hard		SVAMP		MultiArith		Pass@1 ↑	Average	Pass@1/# L ↑
	Pass@1 ↑	# L ↓	Pass@1 ↑	# L ↓	Pass@1 ↑	# L ↓	Pass@1 ↑	# L ↓		# L ↓	
<i>LLaMA-3.2-1B-Instruct</i>											
COT-SFT	49.4±.6	70.1±.2	12.1±.2	78.4±.1	60.5±.4	37.2±.1	93.2±.5	38.5±.1	53.8	56.1	0.96
iCoT	19.8±.2	0.00±.0	3.87±.2	0.00±.0	36.4±.5	0.00±.0	38.2±.7	0.00±.0	24.6	0.00	-
COCONUT	23.1±.3	6.00±.0	5.49±.3	6.00±.0	40.7±.7	6.00±.0	41.1±.2	6.00±.0	27.6	6.00	4.60
CODI	13.3±.6	6.00±.0	2.97±.2	6.00±.0	21.7±.7	6.00±.0	19.2±.8	6.00±.0	14.3	6.00	2.38
CoLaR-5	26.8±.2	5.57±.0	5.87±.1	6.53±.0	48.4±.5	2.95±.0	86.4±.4	3.21±.0	41.9	4.57	9.17
CoLaR-2	40.1±.2	12.7±.0	9.08±.0	14.0±.0	54.9±.2	6.11±.0	91.3±.1	7.35±.0	48.8	10.0	4.87
Latent-SFT(4)	46.8±.1	6.81±.0	11.2±.2	7.77±.2	57.6±.3	3.60±.1	94.0±.2	4.15±.0	52.4	5.58	9.39
- w/o Gumbel	44.6±.0	6.64±.0	10.2±.0	7.68±.0	55.6±.1	3.33±.0	93.6±.0	4.11±.0	51.0	5.44	9.38
- Hidden State	36.7±.4	6.61±.1	8.41±.2	7.77±.0	44.4±.2	3.22±.0	90.3±.1	4.06±.0	45.0	5.42	8.30
- w/o LTIM	38.8±.1	6.59±.1	8.67±.1	8.01±.1	48.9±.1	3.34±.1	90.8±.0	4.36±.0	46.8	5.58	8.39
- w/o LTSuM	42.6±.0	6.70±.0	9.98±.1	7.60±.0	53.9±.1	3.36±.0	92.6±.0	4.27±.0	49.8	5.48	9.08
Latent-SFT(2)	52.4±.2	12.8±.0	12.6±.4	15.0±.1	61.3±.4	6.21±.1	96.8±.3	7.19±.2	55.8 ^{+2.0}	10.3 ^{5.4×}	5.42
- w/o Gumbel	50.4±.1	12.4±.0	10.9±.0	13.9±.0	57.8±.1	5.98±.0	93.8±.0	7.21±.0	53.2	9.87	5.39
- Hidden State	41.1±.3	12.8±.0	9.01±.1	14.2±.1	49.6±.2	6.17±.0	91.1±.0	7.29±.0	47.7	10.1	4.72
- w/o LTIM	41.3±.1	12.9±.1	9.32±.0	14.0±.0	53.6±.1	6.25±.1	91.5±.0	7.96±.0	48.9	10.3	4.76
- w/o LTSuM	49.0±.0	12.1±.0	10.3±.0	13.7±.0	56.5±.2	6.22±.0	93.0±.0	7.23±.0	52.2	9.81	5.32
<i>LLaMA-3.2-3B-Instruct</i>											
COT-SFT	61.1±.3	70.5±.2	18.8±.2	80.6±.1	70.9±.2	34.2±.1	98.3±.5	37.4±.1	62.3	55.7	1.12
Latent-SFT(2)	62.4±.1	12.5±.3	18.3±.3	14.6±.2	71.3±.7	5.94±.1	98.7±.4	7.15±.1	62.7 ^{+0.4}	10.0 ^{5.5×}	6.24

Table 1: Experimental results on four low-difficulty datasets. The number in parentheses represents the compression ratio r . “ ” marks the best result, and “ ” marks the second best. - Hidden State refers to removing the Representation Alignment step and directly using the last-layer hidden state as the representation of latent tokens. - w/o Gumbel indicates removing the stochastic Gumbel noise injection, effectively performing deterministic optimization on static soft labels. - w/o LTIM and - w/o LTSuM denote removing the Latent Token Induction Mask and Latent Token Supervision Mask, respectively, reducing the model to a standard autoregressive attention mechanism.

Model	MATH-500			AIME24		
	Pass@1 ↑	# L ↓	Pass@1/# L × 100 ↑	Pass@1 ↑	# L ↓	Pass@1/# L × 100 ↑
COT-SFT	73.8±1.32	2766±319	2.67	17.7±1.23	4016±343	0.44
Latent-SFT (Hidden State)	67.8±0.94	698.1±22.1	9.71	7.78±2.08	1509±114	0.51
Latent-SFT (<i>Superposition Embedding</i>)	79.8±1.35 ^{+6.0}	759.3±36.9 ^{3.6×}	10.51	19.2±1.60 ^{+1.5}	1484±420 ^{2.7×}	1.29

Table 2: Experimental results of DeepSeek-R1-Distill-Qwen-7B on high-difficulty datasets. We primarily examine the performance gap between *Superposition Embedding* and hidden state representations. To ensure a fair comparison, we set the maximum reasoning chain length to 4,096 tokens for all models.

ing explicit COT-SFT on nearly all datasets. This suggests that explicit reasoning does not represent the upper bound of latent reasoning: with an appropriate definition and learning strategy for latent tokens, latent reasoning can exceed explicit reasoning in accuracy. A real reasoning example is provided in Figure 10 (Appendix).

Performance on High-Difficulty Reasoning Tasks. Table 2 shows that the same trend extends to competition-level math on DeepSeek-R1-Distill-Qwen-7B: Latent-SFT with *Superposition Embedding* surpasses COT-SFT on Pass@1 while shortening chains by 2.7–3.6 $\times$, including on AIME24 where prior latent methods often collapse. It also clearly outperforms the hidden-state variant, indicating that vocab-space alignment remains important at larger scale.

Generalization Beyond Mathematics. To test whether the framework is tied to arithmetic, we also

evaluate logical (ProsQA) and commonsense (CommonsenseQA) reasoning under the COCONUT task-specific setting, with a matched LLaMA-3.2-1B-Instruct backbone and $r = 4$. Latent-SFT(4) reaches 97.8% accuracy with 7.8 generated tokens on ProsQA, versus 82.5% and 56.6 tokens for COT-SFT, and 71.3% with 9.4 tokens on CommonsenseQA, versus 52.9% and 75.5 tokens. Both gains exceed those observed on the mathematical benchmarks (Appendix J).

Wall-Clock and FLOP Efficiency. Shorter chains translate into real inference savings. On the full GSM8k-Aug test set with a matched LLaMA-3.2-1B backbone and decoding stack, Latent-SFT(2) is $4.04 \pm 0.07\times$ faster in wall-clock time than COT-SFT at essentially unchanged peak memory, and needs 36.5% fewer analytical FLOPs per example. The superposition product $\mathbf{z} = \mathbf{E}\alpha$ accounts for only 2.9% of that budget, so vocab-




Figure 4: ECR@K analysis of Latent-SFT on the GSM8k-Aug test set. Each curve represents the per-sample ECR@K, with samples sorted along the x-axis by their ECR@300 values. The dashed line indicates the baseline for uncompressed single-chain reasoning. Latent-SFT consistently surpasses this baseline, with higher training compression ratios yielding greater information compression. This demonstrates that Latent-SFT effectively compresses single explicit reasoning paths during latent reasoning.

space alignment is not an inference bottleneck (Appendix F).

7.3 Ablation Study

Necessity of Masking Mechanisms. LTIM enforces causal independence to stabilize training, while LTSuM ensures latent tokens retain sufficient semantic information for decoding. As shown in Table 1, removing either mechanism leads to a significant accuracy drop ($\sim 5\%$), confirming their critical role in maintaining the structural integrity of the reasoning chain.

Latent Representation and Stochastic Optimization. We further validate our specific design choices. First, replacing *Superposition Embedding* with raw hidden states degrades performance, demonstrating the importance of aligning latent tokens with the intrinsic vocabulary manifold. Second, ablating Gumbel-noise injection (w/o Gumbel) leads to a noticeable decline, confirming that stochasticity acts as a vital regularizer (Theorem 1) to prevent overfitting and enhance generalization.




Figure 5: N_{eff} values of Latent-SFT on the Multi-Chain GSM8k-Aug dataset. The legend reports the mean values of both N_{eff} and the Top-2 Score. The dashed line indicates the empirical COT-SFT baseline under the original metric. Overall, Latent-SFT models across various compression ratios consistently exceed this baseline.

8 Essence of Latent Reasoning

A central premise of Latent-SFT is that latent reasoning should go beyond single-chain compression and support *reasoning superposition*. We therefore examine whether the learned latent process preserves information from a single explicit chain while also capturing multiple reasoning paths, using two complementary hypotheses.

Hypothesis 1. Latent Reasoning as the Compression of a Single Reasoning Chain. To quantify how much explicit information is preserved at each step of latent reasoning, we introduce the **Top-K Effective Compression Ratio (ECR@K)**. This metric measures the proportion of explicit chain tokens that fall within the Top-K of the probability distribution at the corresponding latent step. Larger values indicate stronger horizontal compression across the reasoning chain, while values greater than one reflect true compression of the explicit chain. The formal definition is provided in the Appendix M.2.

The experimental results are presented in Figure 4. Across models trained with different compression ratios r and varying K values, the ECR scores exceed 1 for almost all samples, quantitatively validating that latent reasoning compresses a single reasoning path. Moreover, comparing Figure 4(a) and (b), doubling the training compression ratio nearly doubles the ECR score during inference, further demonstrating the effectiveness of our training strategy.

Hypothesis 2. Latent Reasoning as the Superposition of Multiple Reasoning Chains. To examine whether latent reasoning corresponds to the superposition of multiple paths, we define the ef-




Figure 6: Instance-level multi-path analysis of Latent-SFT(2) on Multi-Chain GSM8k-Aug. Five verified paths use distinct computations, while N_{eff} exceeds the empirical COT-SFT baseline at every latent step.

fective global parallelism (N_{eff}). A larger value reflects a higher degree of parallelism. In addition, we introduce the Top-2 Score, defined as the ratio of the probabilities of the two most likely paths. Formal definitions are provided in the Appendix M.3. To assess the parallel reasoning capability of latent models, we construct multiple alternative reasoning chains corresponding to a single explicit reasoning path within the GSM8k-Aug dataset (Multi-Chain GSM8k-Aug dataset in Appendix L).

Figure 5 presents the analysis results. Across models with different compression ratios, most samples exceed the empirical COT-SFT baseline, indicating that latent reasoning supports more candidate paths than explicit SFT under the same metric. On average, evidence is distributed across roughly 3–4 reasoning paths rather than concentrated on a single path. The average Top-2 ratio exceeds 0.6, suggesting strong “two-way competition” and parallel reasoning. Furthermore, as the compression ratio increases, the degree of parallelism also rises. We provide a latent-to-latent matching analysis with a random baseline in Appendix M.4, and a causal intervention study that truncates the latent chain in Appendix K.

Figure 6 makes this behavior concrete at the

instance level. For a problem admitting five manually verified derivations that reach the same answer through genuinely different intermediate computations, the step-level heatmap shows that the inferred latent trajectory retains simultaneous support for several candidates, with N_{eff} exceeding the empirical COT-SFT baseline at every latent step. Because these candidates differ in their intermediate quantities rather than only in surface form, the observed parallel support is not explained by shared syntax tokens alone.

9 Conclusion

In this work, we introduce Latent-SFT, a unified framework that addresses latent reasoning at the token, chain, and learning levels through Latent-Vocab, Latent-Chain, and Latent-Optim. By doing so, we move latent reasoning beyond an unconstrained black box toward a structured process grounded in the vocabulary interface. Empirical evaluations across six mathematical benchmarks show that Latent-SFT consistently outperforms explicit SFT while reducing reasoning length by $2.7\times$ – $5.5\times$. Crucially, our analysis confirms that the learned latent process captures multiple parallel trajectories rather than merely compressing a single explicit chain.

Limitations

Our study has several limitations. First, due to computational constraints, high-difficulty training uses a filtered Open-R1 subset with reasoning chains shorter than 4K tokens, which may underestimate the method’s potential on longer competition-level problems. Second, our evaluation covers mathematical, logical, and commonsense reasoning, but not code generation, tool use, or agentic workflows; whether latent superposition transfers to these settings remains open. Third, latent chains are constructed with fixed-length segmentation and a multi-stage training pipeline, which adds implementation complexity compared with standard COT-SFT. Finally, the learned latent tokens are not directly human-interpretable, limiting fine-grained analysis of intermediate reasoning steps.

Ethical Considerations

Latent-SFT is intended to improve the efficiency of reasoning models, but it also introduces potential risks. Because latent reasoning steps are less directly interpretable than natural-language CoT, failures may be harder to diagnose or audit. In addition, our Multi-Chain analysis uses GPT-5 to generate candidate alternative reasoning chains; although these candidates are manually filtered, generated data may still reflect model biases or subtle reasoning errors. Finally, efficiency improvements in reasoning systems could lower the cost of both beneficial and harmful automated problem solving, so deployment should include task-appropriate monitoring and safeguards.

Acknowledgments

This work was supported in part by the National Natural Science Foundation of China under Grant Nos. 62676390 and 62276248, by the Beijing Nova Program under Grant No. 20250484765, by the Strategic Priority Research Program of the Chinese Academy of Sciences under Grant No. XDB0680302, by the Key Research and Development Program of Xinjiang Uyghur Autonomous Region under Grant No. 2024B03026, by the Director’s Fund Project of State Key Laboratory of AI Safety, and by the Youth Innovation Promotion Association of the Chinese Academy of Sciences under Grant No. 2023111.

References

- Daman Arora and Andrea Zanette. 2025. [Training language models to reason efficiently](#). *CoRR*, abs/2502.04463.
- Natasha Butt, Ariel Kwiatkowski, Ismail Labiad, Julia Kempe, and Yann Ollivier. 2025. [Soft tokens, hard truths](#). *CoRR*, abs/2509.19170.
- Linbo Cao and Jinman Zhao. 2025. [Pretraining on the test set is no longer all you need: A debate-driven approach to qa benchmarks](#). *Preprint*, arXiv:2507.17747.
- Jeffrey Cheng and Benjamin Van Durme. 2024. [Compressed chain of thought: Efficient reasoning through dense representations](#). *CoRR*, abs/2412.13171.
- DeepSeek-AI. 2025. [Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning](#). *CoRR*, abs/2501.12948.
- Jingcheng Deng, Zhongtao Jiang, Liang Pang, Zihao Wei, Liwei Chen, Kun Xu, Yang Song, Huawei Shen, and Xueqi Cheng. 2025a. [Following the autoregressive nature of LLM embeddings via compression and alignment](#). In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing, EMNLP 2025, Suzhou, China, November 4-9, 2025*, pages 12661–12677. Association for Computational Linguistics.
- Jingcheng Deng, Zihao Wei, Liang Pang, Hanxing Ding, Huawei Shen, and Xueqi Cheng. 2025b. [Everything is editable: Extend knowledge editing to unstructured data in large language models](#). In *The Thirteenth International Conference on Learning Representations, ICLR 2025, Singapore, April 24-28, 2025*. OpenReview.net.
- Yuntian Deng, Yejin Choi, and Stuart M. Shieber. 2024. [From explicit cot to implicit cot: Learning to internalize cot step by step](#). *CoRR*, abs/2405.14838.
- Abhimanyu Dubey et al. 2024. [The llama 3 herd of models](#). *CoRR*, abs/2407.21783.
- Luyu Gao, Aman Madaan, Shuyan Zhou, Uri Alon, Pengfei Liu, Yiming Yang, Jamie Callan, and Graham Neubig. 2023. [PAL: program-aided language models](#). In *International Conference on Machine Learning, ICML 2023, 23-29 July 2023, Honolulu, Hawaii, USA*, volume 202 of *Proceedings of Machine Learning Research*, pages 10764–10799. PMLR.
- 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](#). *CoRR*, abs/2412.06769.
- 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](#). In *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks 1, NeurIPS Datasets and Benchmarks 2021, December 2021, virtual*.

- Hugging Face. 2025. [Open r1: A fully open reproduction of deepseek-r1](#).
- Eric Jang, Shixiang Gu, and Ben Poole. 2017. [Categorical reparameterization with gumbel-softmax](#). In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net.
- Arkil Patel, Satwik Bhattamishra, and Navin Goyal. 2021. [Are NLP models really able to solve simple math word problems?](#) In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2021, Online, June 6-11, 2021, pages 2080–2094*. Association for Computational Linguistics.
- Subhro Roy and Dan Roth. 2015. [Solving general arithmetic word problems](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing, EMNLP 2015, Lisbon, Portugal, September 17-21, 2015, pages 1743–1752*. The Association for Computational Linguistics.
- 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](#). *CoRR*, abs/2502.21074.
- Alon Talmor, Jonathan Herzig, Nicholas Lourie, and Jonathan Berant. 2019. [Commonsenseqa: A question answering challenge targeting commonsense knowledge](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pages 4149–4158. Association for Computational Linguistics.
- Wenhui Tan, Jiaze Li, Jianzhong Ju, Zhenbo Luo, Jian Luan, and Ruihua Song. 2025. [Think silently, think fast: Dynamic latent compression of LLM reasoning chains](#). *CoRR*, abs/2505.16552.
- Qwen Team. 2025. [Qwq-32b: Embracing the power of reinforcement learning](#).
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed H. Chi, Quoc V. Le, and Denny Zhou. 2022. [Chain-of-thought prompting elicits reasoning in large language models](#). In *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022*.
- Zihao Wei, Jingcheng Deng, Liang Pang, Hanxing Ding, Huawei Shen, and Xueqi Cheng. 2025. [Mlake: Multilingual knowledge editing benchmark for large language models](#). In *Proceedings of the 31st International Conference on Computational Linguistics, COLING 2025, Abu Dhabi, UAE, January 19-24, 2025, pages 4457–4473*. Association for Computational Linguistics.
- Zihao Wei, Liang Pang, Jiahao Liu, Wenjie Shi, Jingcheng Deng, Shicheng Xu, Zenghao Duan, Fei Sun, Huawei Shen, and Xueqi Cheng. 2026. [The evolution of thought: Tracking llm overthinking via reasoning dynamics analysis](#). *Preprint*, arXiv:2508.17627.
- Haoyi Wu, Zhihao Teng, and Kewei Tu. 2025a. [Parallel continuous chain-of-thought with jacobi iteration](#). *CoRR*, abs/2506.18582.
- Junhong Wu, Jinliang Lu, Zixuan Ren, Gangqiang Hu, Zhi Wu, Dai Dai, and Hua Wu. 2025b. [Llms are single-threaded reasoners: Demystifying the working mechanism of soft thinking](#). *Preprint*, arXiv:2508.03440.
- Heming Xia, Yongqi Li, Chak Tou Leong, Wenjie Wang, and Wenjie Li. 2025. [Tokenskip: Controllable chain-of-thought compression in llms](#). *CoRR*, abs/2502.12067.
- Shicheng Xu, Liang Pang, Huawei Shen, Xueqi Cheng, and Tat-Seng Chua. 2024. [Search-in-the-chain: Interactively enhancing large language models with search for knowledge-intensive tasks](#). In *Proceedings of the ACM on Web Conference 2024, WWW 2024, Singapore, May 13-17, 2024, pages 1362–1373*. ACM.
- Shicheng Xu, Liang Pang, Yunchang Zhu, Jia Gu, Zihao Wei, Jingcheng Deng, Feiyang Pan, Huawei Shen, and Xueqi Cheng. 2026. [RLKD: distilling LLMs’ reasoning via reinforcement learning](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 40(40):34151–34159.
- Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Tom Griffiths, Yuan Cao, and Karthik Narasimhan. 2023. [Tree of thoughts: Deliberate problem solving with large language models](#). In *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023*.
- Boyi Zeng, Shixiang Song, Siyuan Huang, Yixuan Wang, He Li, Ziwei He, Xinbing Wang, Zhiyu Li, and Zhouhan Lin. 2025. [Pretraining language models to ponder in continuous space](#). *CoRR*, abs/2505.20674.
- Qiyuan Zhang, Fuyuan Lyu, Zexu Sun, Lei Wang, Weixu Zhang, Zhihan Guo, Yufei Wang, Irwin King, Xue Liu, and Chen Ma. 2025a. [What, how, where, and how well? A survey on test-time scaling in large language models](#). *CoRR*, abs/2503.24235.
- Zhen Zhang, Xuehai He, Weixiang Yan, Ao Shen, Chenyang Zhao, Shuohang Wang, Yelong Shen, and Xin Eric Wang. 2025b. [Soft thinking: Unlocking the reasoning potential of llms in continuous concept space](#). *CoRR*, abs/2505.15778.
- Jinman Zhao and Xueyan Zhang. 2024. [Large language model is not a \(multilingual\) compositional relation reasoner](#). In *First Conference on Language Modeling*.

Ruijie Zhu, Tianhao Peng, Tianhao Cheng, Xingwei Qu, Jinfa Huang, Dawei Zhu, Hao Wang, Kaiwen Xue, Xuanliang Zhang, Yong Shan, Tianle Cai, Taylor Kergan, Assel Kembay, Andrew Smith, Chenghua Lin, Binh Nguyen, Yuqi Pan, Yuhong Chou, Zefan Cai, and 14 others. 2025. [A survey on latent reasoning](#). *CoRR*, abs/2507.06203.

Yufan Zhuang, Liyuan Liu, Chandan Singh, Jingbo Shang, and Jianfeng Gao. 2025. [Text generation beyond discrete token sampling](#). *CoRR*, abs/2505.14827.

A Use of LLMs

We employ GPT-5 for two purposes. First, we use it to polish the language of our manuscript: we draft the text, ask GPT-5 to refine it, and then verify its logical and semantic consistency before incorporating the revision. Second, we use GPT-5 to generate preliminary candidate alternative reasoning chains for the Multi-Chain GSM8k-Aug analysis. These generated candidates are not used directly; they are manually inspected, filtered for format and logical correctness, and further filtered for diversity as described in Appendix L.

B Code and Data Availability

The models and datasets of our method are available at <https://github.com/DJC-GO-SOLO/Latent-SFT>.

C Distribution Gap after COCONUT Optimization

To examine whether latent-reasoning training reduces the mismatch between hidden states and token embeddings, we evaluate the released `agurung/coconut-qwen2.5-7b` checkpoint under the same measurement protocol as Section 4. As shown in Table 3, COCONUT optimization does not reduce the distributional discrepancy in this checkpoint; the hidden-state norm and FID both increase compared with the pretrained Qwen-2.5-7B model.

Metric	Qwen-2.5-7B	coconut-qwen2.5-7b
hidden L2	110.27	134.52
embed L2	0.08	0.09
FID	14268.12	17421.55

Table 3: Distribution gap between latent hidden states and the token-embedding space before and after COCONUT optimization.

D Structured Anisotropy of Vocab-Space

We analyze the singular value spectra of token embedding matrices across various LLMs to examine the geometry of vocab-space. Figure 8 shows that spectral energy is distributed unevenly across directions, indicating structured anisotropy rather than an unconstrained isotropic vector space.

E Training Computational Costs

We provide a detailed breakdown of the computational time required for both the **Latent-Chain**




Figure 7: Visualization of last-layer hidden states and token embeddings of the DeepSeek-R1-Distill-Qwen-7B model on the MATH-500 dataset.




Figure 8: Decay curves of singular values for embedding matrices in various LLMs. The legend indicates the effective rank of each embedding matrix.

Construction and Latent-Optim phases. All models were trained on $8 \times$ NVIDIA A100 GPUs.

The training process for Latent-SFT is divided into two primary steps. The first step, **Latent-Chain Construction**, involves three sub-phases: (1) latent-token encoder training, (2) LLM training, and (3) joint fine-tuning. The second step, **Latent-Optim**, involves training the LLM using the proposed stochastic Gumbel-Softmax strategy.

Table 4 summarizes the training duration for each model size. Despite the multi-stage nature of our framework, the total training cost remains manageable.

F Inference Efficiency: Wall-Clock Time and FLOPs

Reasoning length (# L) counts how many steps a model emits, but it does not by itself establish that latent reasoning is cheaper to run. We therefore benchmark end-to-end inference on the complete 1,319-example GSM8k-Aug test set using 8 NVIDIA H20 GPUs. COT-SFT and Latent-SFT(2) share the same LLaMA-3.2-1B backbone, prompts, bf16 precision, FlashAttention-2 backend, per-GPU batch size, and sampling configura-

tion, and each method runs under its own released evaluation cap. We repeat the benchmark three times and report the mean $\pm$ sample standard deviation.

Table 5 summarizes the results. Latent-SFT(2) is $4.04 \pm 0.07 \times$ faster in wall-clock time and attains correspondingly higher throughput. Tail latency improves slightly more than median latency ($4.34 \times$ at P95 versus $3.99 \times$ at P50), because long explicit chains dominate the slowest examples. Peak GPU memory is effectively unchanged, and the method-native generation rate is comparable across the two systems, confirming that the speedup comes from emitting far fewer steps rather than from a faster per-step kernel.

FLOP Accounting. We compute analytical FLOPs under a standard single-prefill implementation with KV-cache reuse, using the actual prompt and generation lengths and including the full-vocabulary superposition cost. Latent-SFT(2) requires 0.2293 TFLOPs per example against 0.3609 for COT-SFT, a 36.5% reduction. The superposition product $\mathbf{z} = \mathbf{E}\alpha$ adds only 0.0067 TFLOPs per example, roughly 2.9% of the Latent-SFT total: it is a single matrix multiplication per latent step, whereas an ordinary decoding step performs the much larger attention Q/K/V/O and MLP projections at every layer. Vocab-space alignment is therefore not an inference bottleneck.

Full-Vocabulary versus Top-K Superposition.

All results in this paper use full-vocabulary superposition, which gives the best task performance. Our implementation also supports restricting α to its Top-K components. Truncating the support discards part of the distribution and costs some accuracy in our experiments, but the resulting bounded

Base Model	Latent-Chain Construction (h)				Latent-Optim (h)	Total (h)
	Latent Enc.	LLM	Joint	Subtotal		
LLaMA-3.2-1B-Instruct	4.0	4.0	4.5	12.5	12.0	24.5
LLaMA-3.2-3B-Instruct	6.5	6.5	8.0	21.0	22.0	43.0
DeepSeek-R1-Distill-Qwen-7B	12.0	12.0	15.0	39.0	50.0	89.0

Table 4: Training time breakdown (in hours) for Latent-SFT across different model scales.

Metric	COT-SFT	Latent-SFT(2)	Change
Pass@1 (%) $\uparrow$	49.4 $\pm$ 0.6	52.4 $\pm$ 0.2	+3.0 pp
Length # L $\downarrow$	70.1 $\pm$ 0.2	12.8 $\pm$ 0.0	5.5 $\times$
Wall-clock (s) $\downarrow$	119.98 $\pm$ 1.26	29.68 $\pm$ 0.20	4.04 $\pm$ 0.07 $\times$
Throughput (ex./s) $\uparrow$	10.99 $\pm$ 0.12	44.44 $\pm$ 0.30	4.04 $\pm$ 0.07 $\times$
P50 latency (s/ex.) $\downarrow$	0.662 $\pm$ 0.002	0.166 $\pm$ 0.001	3.99 $\pm$ 0.03 $\times$
P95 latency (s/ex.) $\downarrow$	1.135 $\pm$ 0.007	0.261 $\pm$ 0.001	4.34 $\pm$ 0.04 $\times$
Peak mem. (GiB/GPU)	2.354	2.356	$\approx$ 1.00 $\times$
FLOPs (TFLOPs/ex.) $\downarrow$	0.3609	0.2293	-36.5%
Gen. units/s	852.91 $\pm$ 8.98	846.31 $\pm$ 5.68	$\approx$ 1.00 $\times$

Table 5: End-to-end inference cost on the full GSM8k-Aug test set (1,319 examples, $8 \times$ NVIDIA H20 GPUs, LLaMA-3.2-1B backbone, three repeats). Latent-SFT(2) is roughly $4\times$ faster in wall-clock time at unchanged peak memory, and requires 36.5% fewer analytical FLOPs.

support makes the latent action space substantially easier to optimize in downstream reinforcement learning, particularly with Group Relative Policy Optimization (GRPO). We therefore treat Top-K superposition as an optional interface rather than the default configuration.

G Theoretical Analysis of Stochastic Latent Optimization

In this section, we analyze the implicit regularization effect of Gumbel-noise injection. Unlike standard Gaussian noise, Gumbel noise has a non-zero mean and fixed variance. We show that, combined with the translation invariance of Softmax, this induces a Hessian-based regularization term in expectation.

Theorem 1 (*Implicit Hessian Regularization via Gumbel-Softmax*). Let $\alpha \in \Delta^{V-1}$ be a static target distribution in the simplex interior (i.e., $\alpha_i > 0$ and $\sum_i \alpha_i = 1$). Consider the stochastic target generated via Eq. (6): $\tilde{\mathbf{p}} = \sigma((\log \alpha + \lambda \mathbf{g})/\tau_g)$, where $\sigma(\cdot)$ is Softmax, $\lambda \geq 0$ is the noise strength, $\tau_g > 0$ is the Gumbel-Softmax temperature, and $\mathbf{g} \in \mathbb{R}^V$ has i.i.d. entries $g_i \sim \text{Gumbel}(0, 1)$. Define the unperturbed baseline

$$\mathbf{p}_{\tau_g} \triangleq \sigma(\tau_g^{-1} \log \alpha). \quad (8)$$

We analyze the expected loss in a neighborhood

of $\mathbf{p}_{\tau_g}$ under the small-perturbation regime (defined below). Then, for a twice-differentiable loss $\mathcal{L}(\theta, \mathbf{p})$ that depends on $\mathbf{p}$ through the forward computation (hence generally non-linear in $\mathbf{p}$), we have:

$$\begin{aligned} \mathbb{E}_{\mathbf{g}}[\mathcal{L}(\theta, \tilde{\mathbf{p}})] &= \mathcal{L}(\theta, \mathbf{p}_{\tau_g}) \\ &\quad + \frac{\lambda^2 \pi^2}{12 \tau_g^2} \text{Tr}(\mathbf{J}_{\sigma}(\mathbf{p}_{\tau_g})^\top \nabla_{\mathbf{p}}^2 \mathcal{L}(\theta, \mathbf{p})|_{\mathbf{p}=\mathbf{p}_{\tau_g}} \\ &\quad \quad \mathbf{J}_{\sigma}(\mathbf{p}_{\tau_g})) \\ &\quad + \mathcal{O}(\|\nabla_{\mathbf{p}} \mathcal{L}\| \mathbb{E}\|\delta\|^2) \\ &\quad + \mathcal{O}(\mathbb{E}\|\delta\|^3), \end{aligned} \quad (9)$$

where $\mathbf{J}_{\sigma}(\mathbf{p}_{\tau_g}) = \text{diag}(\mathbf{p}_{\tau_g}) - \mathbf{p}_{\tau_g} \mathbf{p}_{\tau_g}^\top$ is the Softmax Jacobian evaluated at $\mathbf{p}_{\tau_g}$, and δ is the effective (zero-mean) perturbation in logit space defined in the proof. In particular, in our experiments with $\lambda = 1$ and $\tau_g = 1$, $\mathbf{p}_{\tau_g} = \sigma(\log \alpha) = \alpha$, so the expansion is centered at the semantic target. This indicates that Gumbel noise induces a curvature penalty that favors flatter neighborhoods around the target distribution, with leading-order scaling λ^2/τ_g^2 .

Proof.

1. Notation and mean removal via translation invariance. Let $\mathbf{l} = \log \alpha$. The stochastic target is

$$\tilde{\mathbf{p}} = \sigma\left(\frac{\mathbf{l} + \lambda \mathbf{g}}{\tau_g}\right). \quad (10)$$

A standard Gumbel variable has mean $\mathbb{E}[g] = \gamma$ and variance $\text{Var}(g) = \pi^2/6$. Decompose $\mathbf{g} = \epsilon + \gamma \mathbf{1}$ where $\mathbb{E}[\epsilon] = \mathbf{0}$ and $\mathbb{E}[\epsilon \epsilon^\top] = \frac{\pi^2}{6} \mathbf{I}$. Using Softmax translation invariance $\sigma(\mathbf{x} + c \mathbf{1}) = \sigma(\mathbf{x})$, we obtain

$$\tilde{\mathbf{p}} = \sigma\left(\frac{\mathbf{l} + \lambda \epsilon}{\tau_g}\right) = \sigma\left(\frac{\mathbf{l}}{\tau_g} + \delta\right), \quad \delta \triangleq \frac{\lambda \epsilon}{\tau_g}. \quad (11)$$

Define the baseline $\mathbf{p}_{\tau_g} \triangleq \sigma(\mathbf{l}/\tau_g)$, so that $\tilde{\mathbf{p}} = \sigma(\mathbf{z}_{\tau_g} + \delta)$ with $\mathbf{z}_{\tau_g} = \mathbf{l}/\tau_g$ and $\tilde{\mathbf{p}}(\mathbf{0}) = \mathbf{p}_{\tau_g}$.

2. Second-order expansions. We assume a small-perturbation regime where $\|\delta\|$ is sufficiently

small such that second-order Taylor expansions are valid and the remainders are controlled by $\mathcal{O}(\|\delta\|^3)$ in expectation.

Softmax expansion. Let $\mathcal{H}_\sigma$ denote the (third-order) Hessian tensor of Softmax with respect to logits. Using tensor contraction notation $\mathcal{H}_\sigma[\delta, \delta] \in \mathbb{R}^V$,

$$\tilde{\mathbf{p}} - \mathbf{p}_{\tau_g} = \mathbf{J}_\sigma(\mathbf{p}_{\tau_g}) \delta + \frac{1}{2} \mathcal{H}_\sigma[\delta, \delta] + \mathcal{O}(\|\delta\|^3). \quad (12)$$

Loss expansion. Expanding $\mathcal{L}(\theta, \mathbf{p})$ around $\mathbf{p} = \mathbf{p}_{\tau_g}$,

$$\begin{aligned} \mathcal{L}(\theta, \tilde{\mathbf{p}}) &= \mathcal{L}(\theta, \mathbf{p}_{\tau_g}) + (\tilde{\mathbf{p}} - \mathbf{p}_{\tau_g})^\top \nabla_{\mathbf{p}} \mathcal{L}(\theta, \mathbf{p})|_{\mathbf{p}=\mathbf{p}_{\tau_g}} \\ &\quad + \frac{1}{2} (\tilde{\mathbf{p}} - \mathbf{p}_{\tau_g})^\top \nabla_{\mathbf{p}}^2 \mathcal{L}(\theta, \mathbf{p})|_{\mathbf{p}=\mathbf{p}_{\tau_g}} \\ &\quad (\tilde{\mathbf{p}} - \mathbf{p}_{\tau_g}) + \mathcal{O}(\|\tilde{\mathbf{p}} - \mathbf{p}_{\tau_g}\|^3). \end{aligned} \quad (13)$$

Substituting the first-order relation $\tilde{\mathbf{p}} - \mathbf{p}_{\tau_g} = \mathbf{J}_\sigma \delta + \mathcal{O}(\|\delta\|^2)$ into the quadratic term yields

$$\begin{aligned} \mathcal{L}(\theta, \tilde{\mathbf{p}}) &= \mathcal{L}(\theta, \mathbf{p}_{\tau_g}) + \left(\mathbf{J}_\sigma \delta + \frac{1}{2} \mathcal{H}_\sigma[\delta, \delta] \right)^\top \\ &\quad \nabla_{\mathbf{p}} \mathcal{L}|_{\mathbf{p}=\mathbf{p}_{\tau_g}} + \frac{1}{2} \delta^\top \mathbf{J}_\sigma^\top \nabla_{\mathbf{p}}^2 \mathcal{L}|_{\mathbf{p}=\mathbf{p}_{\tau_g}} \\ &\quad \mathbf{J}_\sigma \delta + \mathcal{O}(\|\delta\|^3), \end{aligned} \quad (14)$$

where $\mathbf{J}_\sigma = \mathbf{J}_\sigma(\mathbf{p}_{\tau_g})$ for brevity.

3. Taking expectations. Since $\mathbb{E}[\delta] = \mathbf{0}$, the linear-in- δ term vanishes: $\mathbb{E}[\mathbf{J}_\sigma \delta] = \mathbf{0}$. The remaining first-order bias arises from Softmax non-linearity:

$$\begin{aligned} &\mathbb{E} \left[\left(\frac{1}{2} \mathcal{H}_\sigma[\delta, \delta] \right)^\top \nabla_{\mathbf{p}} \mathcal{L}|_{\mathbf{p}=\mathbf{p}_{\tau_g}} \right] \\ &= \mathcal{O}(\|\nabla_{\mathbf{p}} \mathcal{L}\| \mathbb{E}[\|\delta\|^2]), \end{aligned} \quad (15)$$

which becomes negligible near a stationary point where $\|\nabla_{\mathbf{p}} \mathcal{L}\| \approx 0$.

For the quadratic term, let $\mathbf{H}_{\tau_g} = \nabla_{\mathbf{p}}^2 \mathcal{L}|_{\mathbf{p}=\mathbf{p}_{\tau_g}}$ and $\mathbf{A}_{\tau_g} = \mathbf{J}_\sigma^\top \mathbf{H}_{\tau_g} \mathbf{J}_\sigma$. Then

$$\begin{aligned} &\mathbb{E} \left[\frac{1}{2} \delta^\top \mathbf{A}_{\tau_g} \delta \right] \\ &= \frac{1}{2} \text{Tr}(\mathbf{A}_{\tau_g} \mathbb{E}[\delta \delta^\top]). \end{aligned} \quad (16)$$

Using $\mathbb{E}[\epsilon \epsilon^\top] = \frac{\pi^2}{6} \mathbf{I}$ and $\delta = \lambda \epsilon / \tau_g$ gives $\mathbb{E}[\delta \delta^\top] = \frac{\lambda^2 \pi^2}{6 \tau_g^2} \mathbf{I}$, hence

$$\begin{aligned} &\mathbb{E} \left[\frac{1}{2} \delta^\top \mathbf{A}_{\tau_g} \delta \right] \\ &= \frac{\lambda^2 \pi^2}{12 \tau_g^2} \text{Tr}(\mathbf{A}_{\tau_g}). \end{aligned} \quad (17)$$

Collecting terms yields Eq. (9), with the remainder $\mathcal{O}(\mathbb{E}[\|\delta\|^3])$ from Taylor truncation. $\square$

H Latent-Chain Construction Training

The strict constraints imposed by LTIM and LT-SuM can destabilize naive SFT optimization. To address this, we adopt an EM-style alternating optimization schedule. Let $\theta = \{\theta_{\text{enc}}, \theta_{\text{llm}}\}$ represent the trainable parameters, where θ_{enc} denotes the latent-token encoder initialized from the base LLM and θ_{llm} denotes the target LLM. The training proceeds in three phases:

1. Latent-Token Encoder Optimization:

Freezing the LLM backbone, we update θ_{enc} to minimize reconstruction loss: $\theta_{\text{enc}}^{(t+1)} \leftarrow \arg \min_{\theta_{\text{enc}}} \mathcal{L}_{\text{sup}}(\theta_{\text{enc}}, \theta_{\text{llm}}^{(t)})$.

2. LLM Optimization:

Freezing the latent-token encoder, we update the LLM backbone to adapt to the latent representations: $\theta_{\text{llm}}^{(t+1)} \leftarrow \arg \min_{\theta_{\text{llm}}} \mathcal{L}_{\text{sup}}(\theta_{\text{enc}}^{(t+1)}, \theta_{\text{llm}})$.

3. Joint Fine-tuning:

Finally, we jointly optimize both components to ensure global convergence.

This progressive strategy stabilizes the training dynamics, allowing the model to gradually adapt to the distributional shift introduced by the latent tokens.

H.1 Fixed-Length Segmentation

In our experiments, $\text{Seg}(\cdot)$ uses a fixed compression ratio r . Given a tokenized explicit reasoning chain $C = (c_1, \dots, c_L)$, we set

$$N = \left\lceil \frac{L}{r} \right\rceil, \quad (18)$$

and define each segment as

$$S_i = (c_{(i-1)r+1}, \dots, c_{\min(ir, L)}), \quad i = 1, \dots, N. \quad (19)$$

We then insert one latent placeholder L_i immediately after each segment S_i , yielding the augmented chain $[S_1, L_1, S_2, L_2, \dots, S_N, L_N]$. The label of each placeholder token is set to -100 so that it is ignored by the standard token-level cross-entropy loss.

I Implementation Details

Evaluation Metrics and Protocol. Our evaluation measures both accuracy (Pass@1) and efficiency (# L, the number of tokens in the reasoning chain). For the low-difficulty datasets, we conducted five runs with different random seeds. For the high-difficulty datasets, MATH-500 was evaluated four times, while AIME24 was evaluated 64 times. We report Pass@1 and # L as the mean $\pm$ 95% confidence interval (CI). To jointly assess accuracy and efficiency, we also report their ratio, Pass@1/# L.

Training and Inference Setup. Since Latent-SFT requires weighting with respect to the vocabulary matrix, we freeze the vocabulary parameters during training. All models are trained using LoRA with rank 64. Following COCONUT, we initialize all methods from the COT-SFT model to accelerate convergence. In practice, we find that a fixed-length segmentation function performs better than a semantics-based segmentation. Therefore, all results reported in this paper are based on fixed-length segmentation functions. For LLaMA-3.2-1B-Instruct and LLaMA-3.2-3B-Instruct, we set the learning rate to 1×10^{-4} for Latent-Chain generation and 4×10^{-4} for Latent-Optim. For DeepSeek-R1-Distill-Qwen-7B, the learning rates are 3×10^{-5} (Latent-Chain) and 1×10^{-5} (Latent-Optim). Following DeepSeek-R1 (DeepSeek-AI, 2025), we use a decoding temperature of 0.6 and top-p of 0.95.

During evaluation, low-difficulty mathematical reasoning datasets were evaluated using the Transformers library with a maximum generation length of 128 tokens. High-difficulty datasets were evaluated using SGLang with a context-length limit of 32,768 tokens and a maximum generated reasoning-chain length of 4,096 tokens. Except for the system-efficiency benchmark in Appendix F, which used 8 NVIDIA H20 GPUs, all training and accuracy-evaluation experiments were conducted on 8 NVIDIA A100 GPUs.

Baseline Fairness Protocol. The iCOT, COCONUT, CODI, and CoLaR numbers in Table 1 are taken from the standardized comparison of Tan et al. (2025), which tunes these baselines under a unified evaluation framework, rather than being re-implemented by us with reduced budgets. Across all methods we control the base model, training and evaluation data, LoRA rank, COT-SFT ini-

Dataset	Method	Accuracy (%) $\uparrow$	Avg. tokens $\downarrow$
ProsQA	COT-SFT	82.5	56.6
	Latent-SFT(4)	97.8	7.8
CommonsenseQA	COT-SFT	52.9	75.5
	Latent-SFT(4)	71.3	9.4

Table 6: Logical (ProsQA) and commonsense (CommonsenseQA) reasoning under the COCONUT task-specific setting, with a matched LLaMA-3.2-1B-Instruct backbone for both methods. Avg. tokens is the average number of generated tokens per example.

tialization quality, and decoding settings. Learning rates and update counts are allowed to differ, because these methods optimize materially different objectives and forcing identical optimization hyperparameters would push some of them into known suboptimal regimes. For the high-difficulty benchmarks, prior latent-reasoning methods neither report MATH-500 or AIME24 results nor release validated checkpoints for these settings, so Table 2 isolates the effect of vocab-space superposition against matched COT-SFT and hidden-state variants instead of reproducing every prior method from scratch.

J Generalization to Logical and Commonsense Reasoning

The six primary benchmarks in Tables 1 and 2 are mathematical, so we additionally test whether Latent-SFT transfers to reasoning that is not arithmetic. Following the task-specific protocol of COCONUT (Hao et al., 2024), we train and evaluate on ProsQA, a planning-style logical deduction benchmark introduced in that work, and on CommonsenseQA (Talmor et al., 2019), a multiple-choice commonsense benchmark. COT-SFT and Latent-SFT use the same LLaMA-3.2-1B-Instruct backbone, the same training data, and the same evaluation protocol; Latent-SFT is trained with compression ratio $r = 4$.

Table 6 reports the results. On ProsQA, Latent-SFT improves accuracy by 15.3 points while reducing the average generated-token count by 86.2% (7.8 versus 56.6 tokens, a compression ratio of 7.26). On CommonsenseQA, it improves accuracy by 18.4 points while reducing generated tokens by 87.5% (9.4 versus 75.5, a ratio of 8.03). Both improvements are considerably larger than those on the mathematical benchmarks in Table 1, and the matched backbone rules out model scale as the source of the difference. This pattern is consistent with COCONUT and related latent-reasoning work,

Dataset	Full	25%	50%	75%
GSM8k-Aug	52.4	46.4	33.9	21.0
MATH-500	79.8	63.2	51.1	40.5

Table 7: Pass@1 after truncating the trailing 25%, 50%, and 75% of the generated latent chain and decoding the answer from the remaining prefix. GSM8k-Aug uses LLaMA-3.2-1B with Latent-SFT(2); MATH-500 uses DeepSeek-R1-Distill-Qwen-7B.

which report the clearest benefits on planning-intensive logical tasks such as ProsQA.

K Latent Intervention Analysis

The ECR@K and N_{eff} analyses are observational: they describe what the latent distributions contain, but they do not establish that the latent chain is causally used when the answer is produced. We therefore run a direct intervention. The trained model first generates its latent chain as usual; we then truncate the trailing x fraction of latent steps and force the model to decode the final answer from the remaining prefix. If latent tokens were inert placeholders whose only role is to buy extra decoding steps, accuracy would be largely insensitive to x .

Table 7 shows the opposite. Accuracy degrades monotonically and steeply with x . On GSM8k-Aug, removing 25%, 50%, and 75% of the latent chain costs 6.0, 18.5, and 31.4 points relative to the intact trajectory. On MATH-500, the corresponding drops are 16.6, 28.7, and 39.3 points; the larger sensitivity is consistent with its longer and less redundant derivations. Together with the distributional evidence of Appendix M.2 and the representation-level evidence of Appendix M.4, this provides functional confirmation that the latent chain carries the intermediate computation rather than merely occupying decoding steps.

L Construction of the Multi-Chain GSM8k-Aug Dataset

We carefully designed prompts and employed GPT-5 to generate preliminary candidate reasoning chains. All generated candidates were manually inspected by the authors, and candidates that failed to follow the required reasoning format or contained incorrect logic were discarded. To further maximize diversity, we retained only reasoning chains with pairwise similarity below 0.9 based on edit distance. An example prompt call is illustrated in

Figure 9. Among the 1,319 data points, 1,310 contained multiple reasoning paths after filtering. The resulting Multi-Chain data are released with the public repository in Appendix B.

M Analytical Metrics

M.1 Notation and Preliminaries

Let $\mathcal{V}$ be the vocabulary, $|\mathcal{V}| = V$. For a given problem, fix a reference explicit chain (from an explicit CoT solution)

$$\mathbf{x} = (x_1, \dots, x_{L_{\text{exp}}}), \quad x_i \in \mathcal{V}. \quad (20)$$

A latent model generates a latent reasoning chain consisting of T latent steps, each with a vocabulary distribution

$$p_t \in \Delta^{V-1} \quad (t = 1, \dots, T), \quad (21)$$

where Δ^{V-1} is the probability simplex. To quantify the extent of explicit information captured by each latent step (whether from a single chain or multiple chains), we introduce:

Assumption 1. During training, if each latent token is constrained to receive information from a fixed number of explicit tokens, then at inference time the amount of explicit information (in terms of token granularity) captured by each latent step is uniform.

This assumption is reasonable, as the model is expected to behave consistently during training and inference. So we align the two sequences by a ratio $r \in \mathbb{N}$: every r explicit tokens are grouped to one latent step. Define the number of aligned steps

$$T' = \min(T, \lceil L_{\text{exp}}/r \rceil), \quad (22)$$

and the explicit token set aligned to latent step t :

$$S_t = \{x_{(t-1)r+1}, \dots, x_{\min(tr, L_{\text{exp}})}\}, \quad (23)$$

$$|S_t| \leq r.$$

Let $\mathcal{T}_K(p_t) \subset \mathcal{V}$ be the set of the K highest-probability tokens under p_t .

M.2 ECR@K (Effective Compression Rate)

To verify whether latent reasoning corresponds to the compression of a single chain, we examine, *on average, how many aligned explicit tokens from the reference chain appear within the Top-K probabilities at each latent step*. Thus, we define

$$\text{ECR@K} = \frac{1}{T'} \sum_{t=1}^{T'} |S_t \cap \mathcal{T}_K(p_t)| \in [0, r]. \quad (24)$$

Instruction:

"Given a GSM8k problem, solution, and answer, identify alternative valid reasoning paths. Enclose each path within `\\boxed{}`. The paths must not contain any natural language, and should follow the same style as the provided example. Ensure the length of the solution and the number of reasoning steps `<>` remain consistent with the standard solution. Each reasoning path must strictly adhere to the problem's logical structure—do not generate unrelated or fabricated paths. If no alternative reasoning path exists, output `\\boxed{False}`."

{"problem": "Janet's ducks lay 16 eggs per day. She eats three for breakfast every morning and bakes muffins for her friends every day with four. She sells the remainder at the farmers' market daily for \$2 per fresh duck egg. How much in dollars does she make every day at the farmers' market?", "solution": "<<16-3-4=9>> <<9*2=18>>", "answer": "18"}"

Response:

"\\boxed{<<<<16-(3+4)=9>> <<9*2=18>>}" **×** Rejected (edit-distance criterion against the solution violated).
 \\boxed{<<3+4=7>> <<(16-7)*2=18>>}"
 \\boxed{<<16-3=13>> <<(13-4)*2=18>>}"
 \\boxed{<<16-4=12>> <<(12-3)*2=18>>}" **×** Rejected (edit-distance criterion against the previous item violated).
 \\boxed{<<2*16=32>> <<32-2*3-2*4=18>>}"

Figure 9: Illustration of the Multi-Chain dataset construction pipeline. We prompt GPT-5 to generate alternative valid reasoning paths for a given GSM8k problem. The response section demonstrates the diversity filtering mechanism: candidate paths are evaluated based on edit distance, and those exhibiting high similarity to the standard solution or previously accepted paths (marked with $\times$) are rejected to ensure significant reasoning diversity.

A value greater than 1 indicates genuine compression of a single reasoning chain, meaning that each latent step, on average, covers more than one token from the reference explicit chain.

M.3 N_{eff} and Top-2 Score

Given M candidate explicit chains $\{x^{(m)}\}_{m=1}^M$ (with lengths L_m), we align each chain to the T latent steps using the same ratio r as above. Let the explicit tokens aligned to step t on chain m be $S_{m,t} \subseteq \mathcal{V}$ (at most r tokens), and let $\mathcal{T}_K(p_t)$ be the Top- K tokens under p_t . Define the per-step *mass* of chain m as

$$\text{mass}_{m,t}^{(K)} = \sum_{x \in S_{m,t} \cap \mathcal{T}_K(p_t)} p_t[x], \quad (25)$$

and aggregate it across steps to obtain a sequence-level evidence score

$$\text{score}_m = \frac{1}{T'} \sum_{t=1}^{T'} \log(\text{mass}_{m,t}^{(K)} + \varepsilon), \quad (26)$$

where $T' = \min(T, \lceil L_m/r \rceil)$, $\varepsilon > 0$ is a small constant, and K controls robustness (setting $K = V$ recovers the full-mass version). We convert $\{\text{score}_m\}$ to a posterior over paths via a temperature-scaled softmax

$$P_m = \frac{\exp(\text{score}_m/\tau)}{\sum_{j=1}^M \exp(\text{score}_j/\tau)} \quad (\tau = 1). \quad (27)$$

The **effective number of supported paths** is

$$N_{\text{eff}} = \exp\left(-\sum_{m=1}^M P_m \log P_m\right), \quad (28)$$

and the **Top-2 score** is the ratio between the two largest posteriors,

$$\text{Top-2} = \frac{P_{(2)}}{P_{(1)}} \quad \text{with} \quad P_{(1)} \geq P_{(2)}. \quad (29)$$

Intuitively, a larger score_m means the model consistently places high probability on chain m along the sequence; a larger N_{eff} and higher Top-2 indicate stronger parallel support across multiple chains rather than collapse to a single path.

When only a single path exists in parallel, $N_{\text{eff}} = 1$. However, different paths may partially overlap in prefixes and syntax tokens, so we use COT-SFT as an empirical baseline instead of a fixed heuristic threshold. On the same Multi-Chain GSM8k-Aug data, COT-SFT obtains $N_{\text{eff}} = 2.01$ with the original metric and $N_{\text{eff}} = 1.66$ after excluding syntax tokens « and » . In addition, to mitigate the influence of tail noise, we set $K = 100$.

M.4 Latent-to-Latent Matching with Random Baseline

The original N_{eff} metric measures support for multiple chains through token-level overlap. While intuitive, it can be affected by fixed positional alignment, shared syntax tokens, or the normalization

in the posterior over paths. To provide a stricter check, we compare inferred latent states directly with the latent states induced by each candidate trajectory, and calibrate the result with a random no-repeat matching baseline.

For each candidate trajectory m , we first encode its explicit segments into reference latent states $\{\mathbf{z}_{m,s}\}_{s=1}^{T_m}$ using the same Latent-Chain construction process. Given an inferred latent chain $\{\hat{\mathbf{z}}_t\}_{t=1}^T$, define $T'_m = \min(T, T_m)$ and compute the latent-to-latent distance matrix

$$D_m[t, s] = \text{dist}(\hat{\mathbf{z}}_t, \mathbf{z}_{m,s}), \quad (30)$$

where $\text{dist}(\cdot, \cdot)$ is the latent-space distance. We then find the best no-repeat alignment

$$\pi_m^* = \arg \min_{\pi \in \Pi_m} \frac{1}{T'_m} \sum_{t=1}^{T'_m} D_m[t, \pi(t)], \quad (31)$$

where Π_m is the set of injective mappings from inferred latent steps to reference segments. The trajectory-level matching score is

$$a_m = -\frac{1}{T'_m} \sum_{t=1}^{T'_m} D_m[t, \pi_m^*(t)]. \quad (32)$$

We convert $\{a_m\}_{m=1}^M$ into a posterior over candidate trajectories and compute the effective number of matched paths:

$$P_m^{\text{match}} = \frac{\exp(a_m/\tau)}{\sum_{j=1}^M \exp(a_j/\tau)},$$

$$N_{\text{eff}} = \exp\left(-\sum_{m=1}^M P_m^{\text{match}} \log P_m^{\text{match}}\right). \quad (33)$$

For calibration, we sample random no-repeat mappings $\pi_{m,b}^{\text{rand}} \in \Pi_m$ and compute the same quantity, yielding $N_{\text{eff,naive}}$. If a high N_{eff} only results from posterior normalization or random path overlap, then $N_{\text{eff}}/N_{\text{eff,naive}}$ should be close to 1. A value greater than 1 indicates that the inferred latent chain matches valid alternative trajectories better than random alignment.

Table 8 reports this comparison. Latent-SFT(2) and Latent-SFT(4) reach N_{eff} of 2.97 and 3.42 against random baselines of 1.86 and 2.51, giving ratios of 1.59 and 1.36. Both ratios are clearly above 1, so the multi-path support measured here is not an artifact of the posterior normalization or of incidental overlap between candidate trajectories.

Model	$\mu(N_{\text{eff}})$	$\mu(N_{\text{eff,naive}})$	$\mu(\frac{N_{\text{eff}}}{N_{\text{eff,naive}}})$
Latent-SFT(2)	2.97	1.86	1.59
Latent-SFT(4)	3.42	2.51	1.36

Table 8: Latent-to-latent matching analysis with a random no-repeat baseline on Multi-Chain GSM8k-Aug.

We also compare the original N_{eff} metric with COT-SFT under the same candidate set. As shown in Table 9, Latent-SFT remains substantially higher than COT-SFT, even after excluding syntax tokens « and » from the counted vocabulary.

Model	Original $\mu(N_{\text{eff}})$	w/o syntax tokens
Latent-SFT(2)	3.60	3.10
Latent-SFT(4)	4.61	3.98
COT-SFT	2.01	1.66

Table 9: Original N_{eff} metric compared with COT-SFT on Multi-Chain GSM8k-Aug.

Question:

"Janet's ducks lay 16 eggs per day. She eats three for breakfast every morning and bakes muffins for her friends every day with four. She sells the remainder at the farmers' market daily for \$2 per fresh duck egg. How much in dollars does she make every day at the farmers' market?"

Reference Answer:

"<<16-3-4=9>> <<9*2=18>>\boxed{18}",

Latent-SFT(2) Response:

"47 <<000=>>9218 <<18\boxed{18}"

Top-1 Tokens in Latent Reasoning

Figure 10: Reasoning example of Latent-SFT(2) on the GSM8k-Aug test set. The model generates a sequence of latent tokens (annotated as "Top-1 Tokens") instead of explicit natural language steps. While these tokens appear as unintelligible numeric combinations to humans, they effectively encode the necessary logical computations, allowing the model to derive the correct answer, 18. Notably, this latent process reduces the reasoning chain length by nearly 50% compared to the reference solution, demonstrating high computational efficiency.

Algorithm 1 Latent-SFT: Latent-Chain Construction, Latent Optimization, and Inference

Require: Question Q , explicit reasoning chain C , final answer A , segmentation function $\text{Seg}(\cdot)$, token embedding matrix $\mathbf{E}$, LM head $\mathbf{W}$, latent-token encoder θ_{enc} , target LLM θ_{llm} , temperatures τ_{enc} and τ_g , Gumbel noise strength λ , maximum latent length T_{max} .

Ensure: A target LLM that performs reasoning through latent superposition tokens.

1: Set $\tau_{\text{enc}} = 1$, $\tau_g = 1$, and $\lambda = 1$ for all reported experiments.

Stage 1: Construct Latent-Chain Targets

2: Split the explicit chain into N segments:

$$(S_1, \dots, S_N) \leftarrow \text{Seg}(C).$$

3: Insert one latent placeholder after each segment:

$$\bar{X} = [Q, \langle \text{think} \rangle, S_1, L_1, S_2, L_2, \dots, S_N, L_N, \langle / \text{think} \rangle, A].$$

4: **for** $i = 1, \dots, N$ **do**

5: **Latent-token induction.** Run the latent-token encoder on $\bar{X}$ under the Latent Token Induction Mask (LTIM), where placeholder L_i attends only to $[Q, \langle \text{think} \rangle, S_1, \dots, S_i]$.

Future segments and latent placeholders are masked.

6: Extract the hidden state at L_i :

$$h_i \leftarrow \theta_{\text{enc}}(L_i; \text{LTIM}).$$

7: **Vocab-space alignment.** Project h_i to the vocabulary simplex and reconstruct the superposition embedding:

$$\alpha_i = \text{softmax}(\mathbf{W}^\top h_i / \tau_{\text{enc}}), \quad z_i = \mathbf{E} \alpha_i.$$

8: **end for**

9: **Semantic sufficiency supervision.** For each i , define the future explicit target

$$Y_i = [S_{i+1}, \dots, S_N, \langle / \text{think} \rangle, A],$$

with $Y_N = [\langle / \text{think} \rangle, A]$.

10: Optimize the construction process under the Latent Token Supervision Mask (LTSuM): when predicting tokens in Y_i , the model uses the mixed prefix

$$[Q, \langle \text{think} \rangle, z_1, S_2, z_2, \dots, S_i, z_i]$$

and the causal prefix inside Y_i . Here S_1 is represented only through z_1 ; future latent tokens $z_{>i}$ and noncausal target tokens are masked.

11: Train Stage 1 progressively:

$$\theta_{\text{enc}} \rightarrow \theta_{\text{llm}} \rightarrow (\theta_{\text{enc}}, \theta_{\text{llm}}),$$

i.e., first optimize the latent-token encoder, then the target LLM, and finally jointly fine-tune both modules.

12: Store the constructed latent-chain target

$$\mathcal{T}(Q, C, A) = \{(\alpha_i, z_i)\}_{i=1}^N.$$

Stage 2: Train the Target LLM with Latent-Optim

13: Remove all explicit intermediate reasoning segments and form the latent training sequence:

$$X_{\text{lat}} = [Q, \langle \text{think} \rangle, \bar{z}_1, \dots, \bar{z}_N, \langle / \text{think} \rangle, A].$$

14: **for** each training update and each latent position i **do**

15: Sample Gumbel noise:

$$g_i \sim \text{Gumbel}(0, 1)^V.$$

16: Construct the stochastic latent distribution:

$$\tilde{p}_i = \text{softmax}\left(\frac{\log \alpha_i + \lambda g_i}{\tau_g}\right).$$

17: Use the same stochastic distribution both as the KL target and as the teacher-forced latent input:

$$\tilde{z}_i = \mathbf{E} \tilde{p}_i.$$

18: **end for**

19: Let $q_i = \text{softmax}(\mathbf{W}^\top h_i^{\text{llm}})$ be the target LLM prediction at latent position i under the teacher-forced stochastic latent prefix. Optimize

$$\mathcal{L}_{\text{auto}} = \mathbb{E}_{g_{1:N}} \left[\frac{1}{N} \sum_{i=1}^N \text{KL}(\tilde{p}_i \parallel q_i) + \frac{1}{|S_{\text{exp}}|} \sum_{t \in S_{\text{exp}}} -\log q_t[y_t] \right],$$

where S_{exp} denotes the ordinary discrete tokens after the latent chain, including $\langle / \text{think} \rangle$ and answer tokens.

Inference

20: Discard the latent-chain construction encoder θ_{enc} . Initialize the reasoning prefix:

$$P_0 = [Q, \langle \text{think} \rangle].$$

21: Set $t \leftarrow 1$ and $\text{terminated} \leftarrow \text{false}$.

22: **while** $t \leq T_{\text{max}}$ **do**

23: The target LLM produces a hidden state h_t from the current prefix P_{t-1} .

24: Convert the hidden state into a vocabulary distribution:

$$\alpha_t = \text{softmax}(\mathbf{W}^\top h_t).$$

25: Decode the next discrete candidate token:

$$y_t \leftarrow \text{Decode}(\alpha_t).$$

26: **if** $y_t = \langle / \text{think} \rangle$ **then**

27: Append the discrete stop token:

$$P_t = [P_{t-1}, \langle / \text{think} \rangle].$$

28: $\text{terminated} \leftarrow \text{true}$.

29: **break**

30: **else**

31: Construct and append the latent superposition token:

$$z_t = \mathbf{E} \alpha_t, \quad P_t = [P_{t-1}, z_t].$$

32: $t \leftarrow t + 1$.

33: **end if**

34: **end while**

35: **if** $\neg \text{terminated}$ **then**

36: **return** an unsuccessful result.

37: **end if**

38: Decode the final answer autoregressively in the standard token space from the resulting prefix.
