

Improving Latent Reasoning in LLMs via Soft Concept Mixing

Kang Wang¹, Xiaocheng Luo¹, Xiangyu Duan^{1*}, Yu Yao¹,
Zaifu Zhan², Min Zhang¹

¹School of Computer Science and Technology, Soochow University, China

²University of Minnesota, Minneapolis, USA

{kwang0822, xcluo, yyao0605}@stu.suda.edu.cn;

{xiangyuduan, minzhang}@suda.edu.cn; zhan8023@umn.edu

Abstract

Unlike human reasoning in abstract conceptual spaces, large language models (LLMs) typically reason by generating discrete tokens at each step, then discard the rich semantic information contained in the output distribution. Recent work such as Soft Thinking has shown that LLMs’ latent reasoning via soft concepts is a promising direction, but LLMs are trained on discrete tokens, leaving the continuous concept space underutilized during training. To address this, we propose Soft Concept Mixing (SCM)¹, a soft concept aware training scheme that directly exposes the model to soft representations during training. Specifically, SCM constructs a soft concept vector by forming a probability-weighted average of embeddings. Then, this vector is mixed into the model’s hidden states, which embody rich contextual information. Finally, the entire latent reasoning process is optimized with Reinforcement Learning (RL). Experiments on six reasoning benchmarks demonstrate that SCM improves the reasoning performance of LLMs, and simultaneously maintains a stable training dynamic.

1 Introduction

Large language models (LLMs) excel in complex reasoning tasks like mathematics, commonsense, and code. Chain-of-Thought (CoT) prompting (Wei et al., 2022; Lobo et al., 2025) and its related works (Guo et al., 2025; Zhang et al., 2023) substantially improve LLM performance by decomposing problems step by step through the generation of intermediate reasoning traces.

However, unlike human cognition which operates in a high-dimensional abstract space to weigh multiple possibilities (Quiroga et al., 2005; Fedorenko and Varley, 2016), standard CoT reasoning constrains the reasoning process strictly to

sequences of discrete tokens (Yao et al., 2023). This constraint may introduce a significant risk of error propagation, where a minor deviation in early reasoning steps—often stemming from sampling randomness or greedy decoding limitations—cascades into an incorrect final conclusion (Turpin et al., 2023). Crucially, the underlying probability distribution often retains the correct reasoning direction within the top-ranked logits, yet this rich semantic information regarding alternative hypotheses is entirely discarded during the discretization process.

Recent studies have explored latent space reasoning (Deng et al., 2025; Xu et al., 2025; Cheng and Van Durme, 2024; Tack et al., 2026; Yang et al., 2024). For instance, Coconut (Hao et al., 2025) uses the last hidden state of the model as the next-step input embedding, thereby achieving a form of “continuous thought” within the latent space. However, its training paradigm relies on a complex, multi-stage training process that requires a large corpus of CoT trajectories for supervision. This approach is not only computationally intensive but also potentially degrades the general capabilities acquired during pre-training. In contrast, a Reinforcement Learning (RL) framework does not require CoT trajectories, enabling the model to explore alternative reasoning paths through outcome-driven optimization.

Another line of work, Soft Thinking (Zhang et al., 2025), leverages the full output probability distribution to construct a concept token, allowing reasoning to proceed in a high-dimensional semantic space. However, as a method applied only at inference time, it introduces a fundamental gap between its continuous reasoning process and the model’s training on discrete tokens. Since the model is never exposed to such continuous representations during its training phase, it struggles to utilize these soft concepts, causing performance instability. This requires the development of a novel framework that empowers the model to effectively

*Corresponding Author

¹Code is available at <https://github.com/wkedk822/SCM>.

explore and internalize continuous concepts, weaving them directly into the fabric of its own rich, contextual understanding during training.

To achieve this, we propose **Soft Concept Mixing (SCM)**, a lightweight framework that integrates soft concept vectors directly into the language model’s training phase. At each reasoning step, SCM not only samples discrete tokens but also leverages the model’s full output probability distribution to generate a soft concept vector via probability-weighted average of embeddings. This vector is mixed into the model’s hidden states, preserving the distribution’s rich latent semantics. SCM additionally introduces a continuous context conditioning pathway between successive steps. This allows subsequent steps to condition on distributional information produced at step t , complementing discrete token generation with continuous context. To optimize the complex policy space that combines both discrete tokens and continuous concepts, we adopt the RL framework based on Group Relative Policy Optimization (GRPO) (Shao et al., 2024), which stably optimizes the latent policy space and effectively guides the model’s internal reasoning process.

Our evaluations on reasoning benchmarks show that SCM improves performance across different model scales. Furthermore, our approach demonstrates high representational stability; Principal Component Analysis (PCA) of hidden states reveals only slight latent shifts, comparable to the GRPO baseline. In summary, our main contributions are as follows:

- We propose SCM, a lightweight framework that integrates soft concepts into hidden states and utilizes reinforcement learning to dynamically optimize these continuous representations.
- Our evaluations on reasoning benchmarks show that SCM consistently improves performance across different model scales and architectures.
- The analysis of computational efficiency and training stability reveals that SCM demonstrates stability during training and incurs marginal inference overhead, offering a highly favorable trade-off for practical deployment.

2 Related work

2.1 Latent reasoning

Training-based Methods. Coconut (Hao et al., 2025) leverages the model’s hidden states as continuous latent vectors for reasoning. By directly utilizing the hidden state from the previous step as the input embedding for the subsequent step, it achieves more flexible and efficient reasoning within the latent space. However, this approach requires a fixed Chain-of-Thought (CoT). Despite its potential, Coconut relies heavily on Ground Truth CoT data and necessitates a complex multi-stage training strategy to internalize the process from explicit to implicit reasoning. This cumbersome training procedure significantly increases computational overhead, thereby limiting its scalability in large-scale scenarios. CoCoMix (Tack et al., 2026), on the other hand, employs Sparse Autoencoders (SAE) to extract "continuous concepts." By interleaving these concepts with the model’s hidden states, it effectively enhances the model’s semantic representation capabilities.

Inference-time Methods. Soft Thinking (Zhang et al., 2025) constructs "concept tokens" via the probability-weighted averaging of token embeddings, enabling the model to explore multiple reasoning paths simultaneously. However, Wu et al. (Wu et al., 2026) observed that despite the inclusion of soft tokens, the distribution tends to rapidly collapse towards the top-1 token, thereby diminishing the benefit of soft tokens in conveying richer information. To address this limitation, the authors proposed Stochastic Soft Thinking, a method that restores diversity by incorporating randomness via Dirichlet or Gumbel-Softmax sampling. Similarly, Mixture of Inputs (MOI) (Zhuang et al., 2025) synthesizes the probability distribution predicted in the previous step with the actually sampled token using Bayesian estimation to generate a novel input representation. This design allows the model to preserve richer internal representations throughout the entire reasoning process.

2.2 Reinforcement Learning

Reinforcement Learning (RL) has become central to optimizing reasoning capabilities, shifting focus from imitating human labels (SFT) to optimizing outcomes. Early approaches like Proximal Policy Optimization (PPO) (Schulman et al., 2017) employ actor-critic architectures that jointly train a state value function V for advantage estimation,




Figure 1: Training Framework for Soft Concept Mixing (SCM).

a design that inevitably incurs high memory overheads.

To address this, Group Relative Policy Optimization (GRPO) (Shao et al., 2024) emerged as a more efficient, critic-free alternative. By estimating baselines directly from group scores rather than relying on a separate value network, GRPO significantly reduces computational overhead. It stabilizes training by directly optimizing the relative advantages of reasoning trajectories within a sampled group, effectively reinforcing reasoning capability.

Hybrid Reasoning Policy Optimization (HRPO) (Yue et al., 2025) introduces a reinforcement learning framework designed to enhance latent reasoning. Unlike traditional methods that rely solely on discrete tokens, HRPO integrates continuous hidden states into the generation process via a learnable gating mechanism. This mechanism gradually incorporates information from prior hidden states into the current token embedding, allowing the model to leverage a "hybrid" representation. However, this architectural addition incurs extra computational overhead due to the gating parameters. Furthermore, achieving optimal performance is non-trivial, as it necessitates extensive experimentation to precisely tune the hyperparameters that govern the balance between discrete and continuous representations during RL training.

These standard RL approaches excel at optimizing reasoning strategies based on outcomes. However, they are often constrained by the rigid dis-

creteness of language generation, which forces the model to make hard token commitments at every step, potentially discarding valuable alternative reasoning paths that could enhance overall solution quality.

3 Methodology

In this section, we introduce Soft Concept Mixing (SCM), a lightweight training framework designed to enhance the latent reasoning capabilities of Large Language Models. The overall process is illustrated in Figure 1.

3.1 Soft Concept Vector Generation

To generate a guidance signal more informative than discrete tokens, we introduce the soft concept vector. At step t , given the output logits $\mathbf{z}_t \in \mathbb{R}^{|V|}$, we compute the probability distribution $p_t = \text{Softmax}(\mathbf{z}_t)$. The soft concept vector $\hat{\mathbf{s}}_t$ is then constructed as the expectation of token embeddings E over this distribution:

$$\hat{\mathbf{s}}_t = \sum_{i=1}^{|V|} p_{t,i} \cdot e(x_i) = \mathbb{E}_{x \sim p_t}[e(x)].$$

In practice, we apply top- k and top- p truncation to suppress low-probability noise and concentrate probability mass on semantically relevant candidates, followed by renormalization. To ensure semantic consistency, these truncation settings are aligned with the decoding configurations used during inference. The resulting vector $\hat{\mathbf{s}}_t$ provides a

continuous summary of the high-probability token embeddings, capturing the dominant distributional concept at the current step.

3.2 Soft Concept Mixing

To effectively inject the rich distributional information carried by the soft concept vector into the model’s latent representation space, we leverage the native self-attention mechanism of the Transformer. We project the soft concept into the Key-Value subspace across all transformer layers and all attention heads, providing it as additional context.

Let $\mathbf{h}_t^{(l)}$ denote the hidden state input of the l -th layer at time step t . In standard self-attention computation, the model first derives the Query vector, along with the Key and Value matrices representing the historical context:

$$\mathbf{q}_t^{(l)} = \mathbf{h}_t^{(l)} \mathbf{W}_Q^{(l)}, \quad \mathbf{K}_{<t}^{(l)}, \quad \mathbf{V}_{<t}^{(l)}$$

where $\mathbf{W}_Q^{(l)}$ is the projection matrix, and $\mathbf{K}_{<t}^{(l)}$ and $\mathbf{V}_{<t}^{(l)}$ denote the sequences of Keys and Values from previous time steps.

To inject the soft concept vector $\hat{\mathbf{s}}e_{t-1}$ generated from the previous time step, we first project it into the current layer’s semantic space to generate the **concept Key** and **concept Value**:

$$\mathbf{k}_{se}^{(l)} = \hat{\mathbf{s}}e_{t-1} \mathbf{W}_K^{(l)}, \quad \mathbf{v}_{se}^{(l)} = \hat{\mathbf{s}}e_{t-1} \mathbf{W}_V^{(l)}$$

Subsequently, we extend the attention context with these concept vectors. The modified Key and Value sequences become:

$$\hat{\mathbf{K}}^{(l)} = [\mathbf{K}_{<t}^{(l)}; \mathbf{k}_{se}^{(l)}], \quad \hat{\mathbf{V}}^{(l)} = [\mathbf{V}_{<t}^{(l)}; \mathbf{v}_{se}^{(l)}]$$

Finally, the model computes the attention output based on this extended context:

$$\mathbf{h}_t'^{(l)} = \text{Attention}(\mathbf{q}_t^{(l)}, \hat{\mathbf{K}}^{(l)}, \hat{\mathbf{V}}^{(l)})$$

Thus, the soft concept $\hat{\mathbf{s}}e_{t-1}$ participates in the reasoning process as a contextual guidance signal. Crucially, this approach ensures geometric consistency between the concept vectors and the model’s original feature space, while empowering the model to flexibly “backtrack” to utilize distribution information within the latent space.

3.3 Optimizing Latent Reasoning via RL

To effectively optimize the reasoning process, we view generation as a discrete token policy augmented with a deterministic continuous conditioning pathway. The discrete policy governs token

selection, while the latent continuous policy generates soft concept vectors that enhance the model’s hidden state, bridging continuous semantic exploration and discrete decision-making.

We formalize the reasoning process as a Markov Decision Process (MDP). At each step t , the state s_t consists of the input $\mathbf{x}$, the history of generated tokens $\mathbf{y}_{<t}$, the current hidden state $\mathbf{h}_t$, and the previous concept vector $\hat{\mathbf{s}}e_{t-1}$. The policy π_θ operates in two sequential phases: 1) **Discrete Action**: The model integrates $\hat{\mathbf{s}}e_{t-1}$ into the attention context to derive the enhanced hidden state $\mathbf{h}_t'$, and subsequently samples the discrete token y_t from the distribution $\mathbf{p}_t = \pi_\theta(\cdot | \mathbf{h}_t')$. 2) **Latent Action**: Using the materialized distribution $\mathbf{p}_t$, a deterministic operation computes the current soft concept vector $\hat{\mathbf{s}}e_t = \mathbb{E}_{v \sim \mathbf{p}_t}[e(v)]$ via a differentiable weighted sum over the vocabulary.

GRPO Training. We optimize π_θ using GRPO (Shao et al., 2024). For each input $\mathbf{x}$, we sample K rollouts. In each rollout, the model applies the concept-mixing operation at every step to obtain the enhanced state $\mathbf{h}_t'$ and then samples $y_t \sim \pi_\theta(\cdot | \mathbf{x}, \mathbf{y}_{<t}, \mathbf{h}_t')$. We compute a total reward $r^{(k)}$ per rollout based on correctness and formatting constraints. Using the group-normalized advantage $A^{(k)} = (r^{(k)} - \mu)/(\sigma + \epsilon)$, our training objective is to minimize:

$$\mathcal{L}(\theta) = -\mathbb{E}_{\mathbf{x}} \left[\frac{1}{K} \sum_{k=1}^K \frac{1}{|y^{(k)}|} \sum_{t=1}^{|y^{(k)}|} \left(\min \left(\rho_t^{(k)} A^{(k)}, \text{clip} \left(\rho_t^{(k)}, 1 - c, 1 + c \right) A^{(k)} \right) - \beta D_{\text{KL}}(\pi_\theta \| \pi_{\text{ref}}) \right) \right]$$

where $\rho_t^{(k)} = \frac{\pi_\theta(y_t^{(k)} | \mathbf{x}, \mathbf{y}_{<t}^{(k)}, \mathbf{h}_t'^{(k)})}{\pi_{\text{old}}(y_t^{(k)} | \mathbf{x}, \mathbf{y}_{<t}^{(k)}, \mathbf{h}_t'^{(k)})}$ denotes the probability ratio, c is the clipping threshold, and π_{ref} is the reference model. This eliminates the need for a separate value function, providing stable gradients to directly reinforce high-reward trajectories.

Continuous Semantic Conditioning. Instead of treating the discarded probability distribution p_t as a byproduct, SCM leverages it to enrich the immediate downstream state. By deterministically injecting the expected semantic vector $\hat{\mathbf{s}}e_t$ into the KV cache, the subsequent state $\mathbf{h}_{t+1}'$ is conditioned on a multi-hypothesis representation rather than a single, hard-committed discrete token. Consequently, when GRPO optimizes the discrete token policy based on outcome rewards, it operates over these semantically enriched states, allowing the model to make more robust decisions grounded in broader latent context.

Model	Method	MATH500	AIME24	GSM8K	GPQA-D	MLLU	BBH	Avg.
DS-R1-Q-7B	CoT	91.75	50.00	90.25	50.44	66.58	45.19	65.70
	Soft Thinking	90.55	48.75	89.40	49.62	66.35	46.24	65.15
	GRPO	93.20	53.33	90.45	51.01	67.03	47.34	67.06
	SCM	94.40	56.25	92.53	52.84	67.91	48.98	68.82
	w/o hidden states	93.00	54.58	90.80	51.07	66.90	45.25	66.93
DS-R1-L-8B	CoT	81.10	44.58	69.35	44.38	71.15	46.33	59.48
	Soft Thinking	77.55	40.42	70.15	42.87	70.60	47.88	58.25
	GRPO	81.43	47.92	70.25	44.89	72.10	48.02	60.77
	SCM	82.60	55.00	71.90	45.58	73.90	48.87	62.98
	w/o hidden states	81.90	50.42	70.60	45.08	71.90	47.16	61.18
DS-R1-Q-1.5B	CoT	80.80	30.00	75.20	33.27	44.38	24.93	48.10
	Soft Thinking	80.60	20.00	77.30	31.82	44.17	26.18	46.68
	GRPO	82.75	33.33	76.95	33.71	45.02	25.57	49.56
	SCM	84.00	36.67	77.79	35.98	46.90	29.44	51.80
	w/o hidden states	80.60	33.33	77.18	33.84	45.80	26.08	49.47
Qwen2.5-7B-INS	CoT	75.00	12.08	89.61	32.58	68.23	59.12	56.10
	Soft Thinking	76.10	13.33	90.96	32.20	68.27	60.04	56.82
	GRPO	76.20	13.33	90.42	32.95	68.86	60.62	57.06
	SCM	77.60	18.75	91.89	33.46	70.11	63.47	59.21
	w/o hidden states	75.80	13.33	90.37	32.32	68.09	60.92	56.81

Table 1: Accuracy of SCM compared to other baselines on six reasoning benchmarks. The GRPO baseline enhances CoT with RL fine-tuning. Best results are in **bold**.

4 Experiments

4.1 Experimental setup

Datasets. We train on two mathematical reasoning datasets—GSM8K (Cobbe et al., 2021) and MATH (Hendrycks et al., 2021b), covering a wide difficulty range from elementary arithmetic to competition-level problems. Furthermore, we evaluate our proposed method and all baselines on six public reasoning benchmarks: GSM8K, MATH500, AIME 2024 (MAA Committees, 2024), GPQA-Diamond (Rein et al., 2024), BIG-Bench Hard (BBH) (Suzgun et al., 2023), and MMLU (Hendrycks et al., 2021a).

Models. We conduct experiments using four open-source LLMs: DeepSeek-R1-Distill-Qwen-7B, DeepSeek-R1-Distill-Llama-8B, DeepSeek-R1-Distill-Qwen-1.5B (Guo et al., 2025), and Qwen2.5-7B-Instruct (Qwen Team, 2025).

RL Setup. We fine-tune our models using the GRPO algorithm, implemented with the Unsloth framework. We apply a trainable LoRA module (Hu et al., 2022) (rank=32, alpha=64), targeting all linear layers in the attention (Q/K/V/O) and MLP (gate/up/down) blocks. The models are trained for one epoch on four NVIDIA A100 GPUs with a global batch size of 64, sampling $K = 8$ rollouts per prompt. We use a learning rate of 5×10^{-6} and a maximum output length of 4096 to-

kens. Comprehensive details regarding the reward formulation and system prompts are provided in Appendix B.

Evaluation Settings. All methods use the same chat templates and CoT prompting, and we report avg@8 accuracy, defined as the average accuracy over 8 independent sampling runs per problem. To assess varying reasoning budgets, extended pass@ k and majority-vote@ k metrics are detailed in Appendix E. For decoding, we use top- $k=30$ sampling with temperature 0.6, top- $p=0.95$, and a maximum generated length of 32,768 tokens. Soft Thinking (Zhang et al., 2025) is evaluated with its official implementation and default hyperparameters, detailed in Appendix C.1.

4.2 Main Results

We evaluated the proposed Soft Concept Mixing (SCM) method on four models, and compared it with traditional Chain-of-Thought (CoT) reasoning, Soft Thinking, and the reinforcement learning baseline of GRPO fine-tuned on CoT. In addition, we conducted an ablation study (denoted as SCM w/o hidden states) where the soft concept vector is directly added to the token embedding of the next step. Detailed results are shown in Table 1. To verify the robustness of these performance improvements, the results of statistical significance

Method	MATH500	AIME 2024	GSM8K	GPQA-Diamond	MMLU	BBH	Avg
CoT Thinking	91.75	50.00	90.25	50.44	66.58	45.19	65.70
Soft Thinking(Gumbel)	93.45	52.08	89.52	50.69	67.35	45.67	66.46
MOI	93.53	52.08	89.04	48.99	66.85	45.98	66.08
HRPO	93.73	55.00	89.88	50.76	66.70	46.34	67.07
Coconut	92.60	43.33	89.25	47.10	66.50	45.27	64.01
SCM (Ours)	94.40	56.25	92.53	52.84	67.91	48.98	68.82

Table 2: Comparison with concurrent methods on DeepSeek-R1-Distill-Qwen-7B. We report the avg@8 accuracy (%) across six benchmarks.

tests are provided in Appendix D.




Figure 2: Training process of GRPO with and without SCM on DeepSeek-R1-Distill-Qwen-1.5B.

First and foremost, SCM demonstrates consistent efficacy across diverse model architectures. On both DeepSeek-R1-Distill-Qwen-7B and DeepSeek-R1-Distill-Llama-8B, SCM achieves superior average accuracies of 68.82% and 62.98%, respectively, outperforming the CoT and Soft Thinking baselines by margins ranging from 3.1 to 4.7 percentage points. Furthermore, SCM exhibits superior performance compared to the strong GRPO reinforcement learning baseline, particularly on challenging benchmarks such as AIME 2024 and BBH. This indicates that by integrating soft concepts directly into the reinforcement learning process, SCM enables the model to actively adapt its internal representations to leverage these continuous signals, thereby achieving more robust reasoning trajectories.

Furthermore, the robustness of SCM is validated on models with limited capacity. On the lightweight DeepSeek-R1-Distill-Qwen-1.5B, SCM still surpasses CoT, Soft Thinking, and GRPO. This stability is also reflected in the training dynamics, where Figure 2 shows SCM exhibits higher and more stable rewards compared to GRPO in later training stages. The method’s applicability extends beyond reasoning-oriented architectures, as SCM also yields consistent gains on an instruction-tuned model Qwen2.5-7B-Instruct.

Finally, our ablation study confirms the necessity of the hidden state fusion component, as its removal degrades performance.

4.3 Comparison with Concurrent Works

We evaluate SCM against recent training-based (Coconut (Hao et al., 2025), HRPO (Yue et al., 2025)) and inference-time (Stochastic Soft Thinking (Wu et al., 2026), MOI (Zhuang et al., 2025)) latent approaches. For fairness, all baselines follow their official implementations and recommended configurations (Appendix C). Table 2 reports the main results. Extended evaluations (Appendix F) show SCM consistently outperforms concurrent works CODI (Shen et al., 2025) and LTPO (Ye et al., 2026).

Compared to existing latent approaches, SCM is a purpose-built mechanism for autonomous latent exploration. Unlike Coconut, which relies on rigid imitative supervision and a computationally expensive multi-stage curriculum, SCM leverages outcome-driven RL to autonomously discover intuition signals without predefined trajectories. Furthermore, SCM circumvents the architectural complexity of HRPO by projecting concepts natively into the Key-Value subspace, preserving geometric consistency without specialized gating modules. Unlike inference-only strategies (e.g., Stochastic Soft Thinking, MOI) that inherently suffer from training-inference mismatch, SCM integrates soft concepts directly during training. Crucially, SCM establishes a continuous semantic pathway, enabling later steps to be dynamically conditioned on the rich distributions generated in earlier steps.

Collectively, SCM maintains robust out-of-domain performance without catastrophic forgetting, exhibiting distinct advantages in reasoning capability, architectural simplicity, and autonomous optimization over existing latent paradigms.




Figure 3: PCA shift of DeepSeek-R1-Distill-Llama-8B across six datasets.

5 Analysis

5.1 Latent Space Shift Analysis

To further investigate the impact of SCM on the internal representation structure of models, we employ Principal Component Analysis (PCA) to visualize and quantitatively evaluate the shift of latent representation centers before and after training across different models (Huan et al., 2026; Zheng et al., 2025; Xu et al., 2026). Specifically, we apply PCA ($n = 2$) to the hidden states of all transformer layers. For each layer l and evaluation dataset $\mathcal{D}$, we compute the centroid of the hidden states $z_{l,\mathcal{D}}$. The latent space deviation is then defined as the Euclidean distance between the centers of the trained and initialized models: $d_{l,\mathcal{D}} = \|z_{l,\mathcal{D}}^{(\text{trained})} - z_{l,\mathcal{D}}^{(\text{init})}\|_2$. The quantitative results in Table 3 are derived by averaging $d_{l,\mathcal{D}}$ across all transformer layers and all datasets within each specific category, providing a robust measure of global representational drift.

Model	Method	Math Datasets	Knowledge Datasets
DS-R1-L-8B	GRPO	0.02	0.07
	SCM	0.06	0.05
DS-R1-Q-1.5B	GRPO	0.04	0.06
	SCM	0.05	0.08

Table 3: PCA shift distances of representation centers across models.

Figure 3 presents PCA distribution comparisons for the same model evaluated across six datasets. We observe that whether trained solely with GRPO or with the additional SCM mechanism, the overall latent representation shows no significant large-scale drift, indicating stable policy optimization.

This stability is quantitatively confirmed across different model architectures in Table 3. Crucially, SCM achieves significant performance gains despite this minimal representational shift. This suggests that the soft concepts facilitate localized adjustments to intermediate representations, rather than inducing large-scale shifts in the overall representation space.

5.2 Ablation on the Added Noise

To determine whether the performance gains of SCM stem from the preservation of multi-hypothesis semantic distributions, or merely from stochastic regularization and dense architectural connectivity within the embedding manifold, we evaluate four information carriers:

- **Gaussian Noise:** A random vector $\epsilon \sim \mathcal{N}(0, \sigma)$ matching SCM’s average magnitude, testing pure stochastic regularization.
- **Random Token Mixture:** An average embedding of uniformly sampled tokens matching the candidate set size, injecting structurally valid but semantically irrelevant noise.
- **Logit Concatenation:** The raw, unnormalized logit vector $z_t \in \mathbb{R}^{|V|}$ concatenated via a linear layer, exploring raw feature transmission.
- **Hard Token Embedding (Sampled):** The exact embedding of the discretely sampled token y_t , providing a contextually accurate but strictly singular semantic signal to isolate the architectural skip-connection effect.

Method	Information Carrier	Integration	GSM8K	AIME 2024
GRPO (Baseline)	None	-	90.45	53.33
Gaussian Noise	Random $\epsilon \sim \mathcal{N}(0, \sigma)$	Add to $\mathbf{h}_t$	90.52	53.33
Random Token Mix	Random Embeddings $\in \mathbb{R}^d$	Add to $\mathbf{h}_t$	89.95	52.50
Logit Concat	Raw Logits $\in \mathbb{R}^{ V }$	Concat to Input	89.80	51.67
Hard Token Embedding	Sampled Embed $\in \mathbb{R}^d$	Add to $\mathbf{h}_t$	91.20	54.17
SCM (Ours)	Soft Concept $\in \mathbb{R}^d$	Add to $\mathbf{h}_t$	92.53	56.25

Table 4: Ablation of information carriers on GSM8K and AIME2024 using DeepSeek-R1-Distill-Qwen-7B. We compare SCM against alternative injection strategies to isolate the impact of multi-hypothesis semantic distributions.

As shown in Table 4, injecting arbitrary continuous signals (Gaussian Noise, Random Token Mixture) or unnormalized dense features (Logit Concatenation) fails to improve upon the vanilla GRPO baseline, and often degrades performance due to semantic conflict or the curse of dimensionality. This confirms that simply adding structural perturbations or raw features is counterproductive; the hidden state requires semantically aligned guidance.

Crucially, while the Hard Token Embedding baseline provides a contextually valid residual signal and slightly stabilizes reasoning, it strictly underperforms SCM. This empirical gap provides support for our core hypothesis: SCM’s superiority is not a byproduct of denser cross-layer connectivity, but stems directly from preserving the alternative semantic trajectories encoded within the continuous probability distribution p_t .

Method	Inference (A100)		Training (1 Epoch)	
	Speed	Overhead	Avg@8	Time
Vanilla GRPO	1,584 tok/s	-	67.06%	~52 h
SCM (Ours)	1,511 tok/s	+4.6%	68.82%	~54 h

Table 5: Efficiency comparison on DeepSeek-R1-Distill-Qwen-7B.

5.3 Computational Efficiency Analysis

We benchmarked the computational cost using DeepSeek-R1-Distill-Qwen-7B on a single NVIDIA A100. As shown in Table 5, SCM introduces a marginal inference overhead of $\sim 4.6\%$ (1,511 vs. 1,584 tokens/s). Leveraging the truncation strategy, the resulting computation is implemented as a highly parallelized matrix-vector product, ensuring that the per-step overhead remains minimal on GPUs. Regarding training, SCM requires slightly more wall-clock time (~ 54 h vs. ~ 52 h per epoch) but yields an Avg@8 improve-

ment of 1.76 percentage points. This demonstrates that SCM offers a highly favorable trade-off, exchanging minor computational cost for substantial improvements in latent reasoning.

6 Conclusion

In this work, we introduced Soft Concept Mixing (SCM), a lightweight training framework that mixes the rich semantic information from the discarded probability distribution back into the model’s latent space, and explicitly preserves these continuous distributional signals via reinforcement learning. Our experiments across multiple benchmarks and architectures demonstrate that this approach consistently improves performance across diverse reasoning benchmarks while maintaining stable training dynamics, offering a computationally efficient strategy for enhancing latent reasoning in LLMs.

Limitations

While SCM demonstrates significant improvements in latent reasoning, several limitations warrant future investigation. First, our method has been validated exclusively within the GRPO framework; it remains to be verified whether other algorithms (e.g., PPO) or standard Supervised Fine-Tuning (SFT) can effectively align soft concept distributions. Second, future research could investigate whether latent reasoning is more concentrated within specific semantic blocks or hierarchical layers, and explore the potential expansion from token-level soft concepts to more comprehensive context-level representations. Finally, although the injected soft concept vectors empirically enhance performance by serving as effective “intuition” signals, these continuous biases are not directly interpretable by humans, and we currently lack qualitative analysis to fully characterize the specific semantic nature of these latent injections.

Acknowledgments

We would like to thank the anonymous reviewers for their helpful comments. This work was supported by the National Natural Science Foundation of China (Grant Nos. 62537001 and 62276179) and the Project Funded by the Priority Academic Program Development of Jiangsu Higher Education Institutions.

References

- Jeffrey Cheng and Benjamin Van Durme. 2024. Compressed chain of thought: Efficient reasoning through dense representations. *arXiv preprint arXiv:2412.13171*.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, et al. 2021. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*.
- Jingcheng Deng, Liang Pang, Zihao Wei, Shichen Xu, Zenghao Duan, Kun Xu, Yang Song, Huawei Shen, and Xueqi Cheng. 2025. Latent reasoning in llms as a vocabulary-space superposition. *arXiv preprint arXiv:2510.15522*.
- Evelina Fedorenko and Rosemary Varley. 2016. Language and thought are not the same thing: evidence from neuroimaging and neurological patients. *Annals of the New York Academy of Sciences*, 1369(1):132–153.
- Luyu Gao, Aman Madaan, Shuyan Zhou, Uri Alon, Pengfei Liu, Yiming Yang, Jamie Callan, and Graham Neubig. 2023. **PAL: Program-aided language models**. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 10764–10799. PMLR.
- Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Peiyi Wang, Qihao Zhu, Runxin Xu, Ruoyu Zhang, Shirong Ma, Xiao Bi, Xiaokang Zhang, Xingkai Yu, Yu Wu, Z. F. Wu, Zhibin Gou, Zhihong Shao, Zhuoshu Li, Ziyi Gao, Aixin Liu, and 175 others. 2025. **Deepseek-r1 incentivizes reasoning in llms through reinforcement learning**. *Nature*, 645(8081):633–638.
- Shibo Hao, Sainbayar Sukhbaatar, DiJia Su, Xian Li, Zhiting Hu, Jason E Weston, and Yuandong Tian. 2025. **Training large language models to reason in a continuous latent space**. In *Second Conference on Language Modeling*.
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2021a. **Measuring massive multitask language understanding**. In *International Conference on Learning Representations*.
- 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**. In *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks*, volume 1.
- Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2022. **LoRA: Low-rank adaptation of large language models**. In *International Conference on Learning Representations*.
- Maggie Ziyu Huan, Yuetai Li, Tianyu Zheng, Xiaoyu Xu, Seungone Kim, Minxin Du, Radha Poovendran, Graham Neubig, and Xiang Yue. 2026. **Does math reasoning improve general LLM capabilities? understanding transferability of LLM reasoning**. In *Forty-third International Conference on Machine Learning*.
- Elita Lobo, Chirag Agarwal, and Himabindu Lakkaraju. 2025. On the impact of fine-tuning on chain-of-thought reasoning. In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 11679–11698.
- MAA Committees. 2024. American invitational mathematics examination (aime). <https://maa.org/maa-invitational-competitions/>. Accessed: 2024-02-10.
- MAA Committees. 2025. American invitational mathematics examination (aime). <https://maa.org/maa-invitational-competitions/>. Accessed: 2025-03-10.
- Shen-Yun Miao, Chao-Chun Liang, and Keh-Yih Su. 2020. A diverse corpus for evaluating and developing english math word problem solvers. In *Proceedings of the 58th annual meeting of the Association for Computational Linguistics*, pages 975–984.
- 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*, pages 2080–2094.
- R Quian Quiroga, Leila Reddy, Gabriel Kreiman, Christof Koch, and Itzhak Fried. 2005. Invariant visual representation by single neurons in the human brain. *Nature*, 435(7045):1102–1107.
- Qwen Team. 2025. **Qwen2.5 technical report**. *Preprint*, arXiv:2412.15115.
- 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*.

- Subhro Roy and Dan Roth. 2015. [Solving general arithmetic word problems](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1743–1752, Lisbon, Portugal. Association for Computational Linguistics.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*.
- Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, YK Li, Yang Wu, et al. 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. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 677–693.
- Xinyuan Song, Keyu Wang, PengXiang Li, Lu Yin, and Shiwei Liu. 2026. Demystifying the roles of llm layers in retrieval, knowledge, and reasoning. In *ICASSP 2026-2026 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 16792–16796. IEEE.
- Mirac Suzgun, Nathan Scales, Nathanael Schärli, Sebastian Gehrmann, Yi Tay, Hyung Won Chung, Aakanksha Chowdhery, Quoc Le, Ed Chi, Denny Zhou, et al. 2023. Challenging big-bench tasks and whether chain-of-thought can solve them. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 13003–13051.
- Jihoon Tack, Jack Lanchantin, Jane Yu, Andrew Cohen, Iliia Kulikov, Janice Lan, Shibo Hao, Yuandong Tian, Jason E Weston, and Xian Li. 2026. [LLM pretraining with continuous concepts](#). In *The Fourteenth International Conference on Learning Representations*.
- Miles Turpin, Julian Michael, Ethan Perez, and Samuel Bowman. 2023. Language models don’t always say what they think: Unfaithful explanations in chain-of-thought prompting. *Advances in Neural Information Processing Systems*, 36:74952–74965.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed 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*, volume 35, pages 24824–24837. Curran Associates, Inc.
- Junhong Wu, Jinliang Lu, Zixuan Ren, Gangqiang Hu, Zhi Wu, Dai Dai, and Hua Wu. 2026. [LLMs are single-threaded reasoners: Demystifying the working mechanism of soft thinking](#). In *The Fourteenth International Conference on Learning Representations*.
- Xiaoyu Xu, Xiang Yue, Yang Liu, Qingqing Ye, Huadi Zheng, Peizhao Hu, Minxin Du, and Haibo Hu. 2026. [Unlearning isn’t deletion: Investigating reversibility of machine unlearning in llms](#). *Preprint*, arXiv:2505.16831.
- Yige Xu, Xu Guo, Zhiwei Zeng, and Chunyan Miao. 2025. [SoftCoT: Soft chain-of-thought for efficient reasoning with LLMs](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 23336–23351, Vienna, Austria. Association for Computational Linguistics.
- Sohee Yang, Elena Gribovskaya, Nora Kassner, Mor Geva, and Sebastian Riedel. 2024. [Do large language models latently perform multi-hop reasoning?](#) In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 10210–10229, Bangkok, Thailand. Association for Computational Linguistics.
- 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. *Advances in neural information processing systems*, 36:11809–11822.
- Wengao Ye, Yan Liang, and Lianlei Shan. 2026. [Thinking on the fly: Test-time reasoning enhancement via latent thought policy optimization](#). In *The Fourteenth International Conference on Learning Representations*.
- 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](#). In *Advances in Neural Information Processing Systems*, volume 38, Main Conference, pages 5501–5530. Curran Associates, Inc.
- Zhen Zhang, Xuehai He, Weixiang Yan, Ao Shen, Chenyang Zhao, and Xin Wang. 2025. [Soft thinking: Unlocking the reasoning potential of LLMs in continuous concept space](#). In *Advances in Neural Information Processing Systems*, volume 38, Main Conference, pages 168990–169012. Curran Associates, Inc.
- Zhuosheng Zhang, Aston Zhang, Mu Li, and Alex Smola. 2023. [Automatic chain of thought prompting in large language models](#). In *The Eleventh International Conference on Learning Representations*.
- Junhao Zheng, Xidi Cai, Shengjie Qiu, and Qianli Ma. 2025. [Spurious forgetting in continual learning of language models](#). In *International Conference on Learning Representations*, volume 2025, pages 67081–67146.
- Yufan Zhuang, Liyuan Liu, Chandan Singh, Jingbo Shang, and Jianfeng Gao. 2025. [Mixture of inputs: Text generation beyond discrete token sampling](#). In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*.

A Benchmark Details

We utilize six benchmarks to evaluate reasoning across diverse domains and difficulty levels:

- **GSM8K** (Cobbe et al., 2021): A dataset of high-quality grade-school mathematics problems, containing 7.5k training and 1.3k test samples.
- **MATH-500** (Hendrycks et al., 2021b): A 500-problem evaluation subset of the MATH test set, comprising challenging competition-level mathematics problems across seven subjects.
- **AIME 2024** (MAA Committees, 2024): A set of advanced competition-level problems from the 2024 American Invitational Mathematics Examination.
- **GPQA-Diamond** (Rein et al., 2024): A graduate-level benchmark requiring deep domain knowledge and complex reasoning, primarily in STEM fields.
- **BIG-Bench Hard (BBH)** (Suzgun et al., 2023): A suite of 23 challenging tasks selected from Big-Bench to evaluate complex reasoning capabilities.
- **MMLU** (Hendrycks et al., 2021a): A massive multitask benchmark covering 57 subjects across STEM, humanities, and social sciences.

B Experimental Details

In this section, we provide the specific configurations used in our experiments, including the reward formulation and system prompts.

B.1 Reward Function Formulation

To guide the reinforcement learning process, we design a composite reward function $r^{(k)}$ that consists of two components: a correctness reward r_{acc} and a format compliance reward r_{fmt} .

- **Correctness Reward (r_{acc}):** This binary term is set to 1 if the model’s final answer matches the ground truth, and 0 otherwise.
- **Format Reward (r_{fmt}):** To ensure the structural integrity of the reasoning trace, we grant +0.25 for the presence of each of the four required structural tags: `<think>`, `</think>`, `<answer>`, and `</answer>`.

B.2 RL Training Prompt

The system prompt used during the Reinforcement Learning (RL) training phase is presented below:

System Prompt for RL Training

Think about the problem and provide your working out. Then put your final answer within `\boxed{}`. The reasoning process and answer are enclosed within `'<think>' '</think>'` and `'<answer>' '</answer>'` tags, respectively, i.e., `<think>{reasoning process} </think> <answer>\boxed{{final answer}} </answer>.`
{Question}

B.3 Mathematical Reasoning Prompts

The evaluation prompt for mathematical reasoning benchmarks (GSM8K, MATH, AIME 2024) is as follows:

System Prompt for Math Benchmarks

Please reason step by step, and put your final answer within `\boxed{}`.
{Question}

B.4 Knowledge-Based Reasoning Prompts

The evaluation prompt for knowledge-based reasoning benchmarks (MMLU, GPQA, BBH) is as follows:

System Prompt for Knowledge-Based Benchmarks

Please solve the following multiple-choice question. Please show your choice in the answer field with only the choice letter, e.g., "answer": "C".
{Question}

C Implementation Details of Concurrent Baselines

To ensure a fair and rigorous comparison, we strictly followed the official training protocols and hyperparameter settings for all concurrent baselines. For training-based methods (Coconut and HRPO), models were trained on the standard splits

of GSM8K and MATH. For inference-time strategies (Soft Thinking, Stochastic Soft Thinking, and MOI), we kept the evaluation setup identical to our SCM runs (using the same chat template, decoding temperature 0.6, maximum output lengths, and avg@8 metric).

C.1 Soft Thinking

We evaluate Soft Thinking (Zhang et al., 2025) using its official codebase. Following the default configuration, the concept tokens are determined using a sampling strategy with temperature 0.6, top- $p = 0.95$, and top- $k = 30$. Regarding the inference termination strategy, we utilize the "Cold Stop" mechanism with an entropy threshold of 0.1 and a length threshold of 256 tokens.

C.2 Coconut Reimplementation

Since the original Coconut (Hao et al., 2025) implementation targeted GPT-2, we adapted the codebase for the 7B model using Low-Rank Adaptation (LoRA) to manage computational costs while maintaining performance. We strictly implemented its signature **multi-stage continuous reasoning curriculum**:

- **Stage 0 (Warm-up):** The model is first trained with standard Supervised Fine-Tuning (SFT) on the target logic chains.
- **Curriculum Stages:** Following the initial warm-up, we progressively increase the number of continuous latent tokens (k) that replace the original discrete reasoning tokens.
- **Training Schedule:** Each stage consists of 3 training epochs. The transition between stages allows the model to gradually adapt from discrete CoT generation to continuous latent reasoning, ensuring the stability of the latent space.

C.3 HRPO Settings

For Hybrid Reasoning Policy Optimization (HRPO) (Yue et al., 2025), we follow the official implementation and hyperparameter configurations described in the original work.

Gating Mechanism. HRPO introduces a learnable gating parameter Λ to balance latent hidden representations and token embeddings. Following the original setup, the initialization is chosen such that the hidden ratio falls within the recommended range [0.95, 0.999]. Sensitivity analysis in

the HRPO study shows that performance remains stable across this interval with only minor variations across knowledge and STEM benchmarks. In our experiments, we adopt $\Lambda = 0.99$, which lies within this effective range and provides competitive performance across tasks.

C.4 Stochastic Soft Thinking

Stochastic Soft Thinking (Wu et al., 2026) utilizes **Gumbel-Softmax sampling** to inject stochastic noise into the token logits prior to normalization, thereby diversifying the mixture of tokens. In our experiments, we set the internal **Gumbel-Softmax distribution hyperparameter** to $\tau = 0.5$, following the optimal configuration reported in the original study to provide a continuous approximation for the categorical distribution.

C.5 Mixture of Inputs (MOI)

Mixture of Inputs (MOI) (Zhuang et al., 2025) blends the discrete sampled token with the discarded token distribution to maintain richer information. The hyperparameter β controls the concentration over the mixing weights, effectively balancing the model’s prior belief (the output distribution) with its observed evidence (the sampled token). We set the concentration parameter to $\beta = 1$.

D Statistical Significance Analysis

In our main experiments (Table 1), we follow the standard practice of reporting the Avg@8 evaluation metric to evaluate the reasoning capabilities of the models. To strictly evaluate the statistical significance of the performance gains achieved by SCM, we conduct additional sampling-based evaluations. Specifically, utilizing the exact same model checkpoints, we evaluate the DS-R1-Q-7B and DS-R1-L-8B backbones across 3 different random seeds during the decoding phase.

Averaged results and standard deviations ($\pm\sigma$) across these 3 independent evaluation runs are presented in Table 6. Furthermore, we perform a paired t-test between the trial outcomes of SCM and the strongest baseline (GRPO). Results demonstrating statistically significant improvements ($p < 0.05$) are highlighted in **bold** with an asterisk (*).

These findings confirm that SCM consistently outperforms standard RL optimization. The analysis demonstrates that SCM delivers reliable and highly significant performance gains on reasoning-intensive benchmarks, exhibiting strong statistical

Model	Method	MATH500	AIME24	GSM8K	GPQA-D	MMLU	BBH	Avg.
DS-R1-Q-7B	GRPO	93.12 $\pm$ 0.42	52.92 $\pm$ 0.59	90.30 $\pm$ 0.32	51.16 $\pm$ 0.48	67.10 $\pm$ 0.22	47.12 $\pm$ 0.65	66.95 $\pm$ 0.28
	SCM (Ours)	94.28* $\pm$ 0.28	55.83* $\pm$ 0.59	92.41* $\pm$ 0.28	52.69* $\pm$ 0.45	67.85 $\pm$ 0.25	48.75 $\pm$ 0.58	68.64* $\pm$ 0.24
DS-R1-L-8B	GRPO	81.25 $\pm$ 0.48	47.64 $\pm$ 0.71	70.10 $\pm$ 0.45	44.95 $\pm$ 0.55	72.05 $\pm$ 0.32	48.15 $\pm$ 0.68	60.69 $\pm$ 0.38
	SCM (Ours)	82.45* $\pm$ 0.35	54.58* $\pm$ 0.59	71.75* $\pm$ 0.32	45.64 $\pm$ 0.42	73.80* $\pm$ 0.28	48.80 $\pm$ 0.55	62.84* $\pm$ 0.30

Table 6: Significance test of SCM. Values are mean $\pm$ SD across 3 evaluation seeds. **Bold*** denotes $p < 0.05$ (paired t-test).

Model	Method	Metric	MATH	AIME 24	GSM8K	GPQA-D	MMLU	BBH	Avg.
DS-R1-Q-1.5B	GRPO	pass@1	81.40	33.33	76.95	33.84	45.10	24.50	49.19
	SCM	pass@1	82.80	36.67	77.79	34.34	45.80	28.20	50.93
	GRPO	maj@8	84.20	36.67	77.10	34.34	46.20	26.20	50.79
	SCM	maj@8	85.20	36.67	78.17	36.36	47.20	30.10	52.28
Qwen2.5-7B-INS	GRPO	pass@1	75.40	13.33	90.67	32.32	68.10	60.56	56.73
	SCM	pass@1	76.20	16.67	91.89	32.83	69.50	62.94	58.34
	GRPO	maj@8	78.60	16.67	92.42	34.34	71.20	64.10	59.56
	SCM	maj@8	80.20	20.00	94.54	35.35	72.80	65.40	61.38

Table 7: Performance Comparison of Pass@1 and Majority-vote@8.

confidence without relying on favorable sampling variance.

E Extended Evaluation: Pass@1 and Majority Voting

To provide a more comprehensive and granular assessment of SCM’s performance under varied inference constraints, we report the pass@1 and majority-vote@8 (maj@8) results for DeepSeek-R1-Distill-Qwen-1.5B and Qwen2.5-7B-Instruct in Table 7. This evaluation examines the robustness of internalized latent reasoning in both single-trial scenarios and high-budget settings where self-consistency is employed.

F Extended Comparison with Latest Baselines

To further substantiate the effectiveness of our proposed approach, we extend our evaluation to include two recently proposed latent reasoning frameworks: CODI (Shen et al., 2025) and LTPO (Ye et al., 2026). Since these works are highly concurrent, we quote their reported performance directly from the original papers (denoted with †) and evaluate SCM under identical evaluation protocols (e.g., adopting the same base models, benchmark datasets, and pass@1 metric) to ensure a fair comparison.

Comparison with LTPO. We compare SCM against LTPO on the Qwen2.5-7B-Instruct

backbone across five mathematical reasoning benchmarks: GSM8K (Cobbe et al., 2021), Math500 (Hendrycks et al., 2021b), ASDiv Aug (Miao et al., 2020), AIME 2024 (MAA Committees, 2024), and AIME 2025 (MAA Committees, 2025). As shown in Table 8, SCM consistently outperforms LTPO. Specifically, our method achieves an overall average improvement of 1.39 percentage points, with notable gains on GSM8K (91.89% vs. 88.17%) and Math500 (76.20% vs. 73.60%).

Comparison with CODI. CODI attempts to compress the chain-of-thought into a continuous space via a self-distillation paradigm. To compare with this approach, we evaluate SCM against CODI using the GPT-2 model. Following CODI’s evaluation protocol, we measure performance on SVAMP (Patel et al., 2021), GSM-Hard (Gao et al., 2023), and MultiArith (Roy and Roth, 2015). As shown in Table 9, SCM outperforms CODI by a clear margin, demonstrating superior continuous reasoning without requiring the complex multi-stage distillation process.

G Ablation Study on Layer-wise Injection

To explore the impact of SCM injection strategies across different architectural depths, as recent research indicates that transformer layers in LLMs exhibit distinct functionalities in reasoning and knowledge retrieval (Song et al., 2026), we conducted an additional layer-wise ablation study on DeepSeek-

Method	GSM8K	Math500	ASDiv Aug	AIME24	AIME25	Avg.
LTPO [†]	88.17	73.60	92.20	16.67	13.33	56.79
SCM (Ours)	91.89	76.20	92.80	16.67	13.33	58.18

Table 8: Comparison with LTPO on Qwen2.5-7B-Instruct. The symbol [†] indicates the pass@1 accuracy reported by (Ye et al., 2026). Best results are in bold.

Method	SVAMP	GSM-Hard	MultiArith	Avg.
CODI [†]	42.90	9.90	92.80	48.53
SCM (Ours)	56.50	11.20	94.30	54.00

Table 9: Comparison with CODI on GPT-2. The symbol [†] indicates the pass@1 accuracy reported by (Shen et al., 2025). Best results are in bold.

R1-Distill-Qwen-1.5B. Specifically, we divided the transformer layers into three groups: Bottom (layers 0–8), Middle (layers 9–18), and Top (layers 19–27), and evaluated the model’s performance when soft concepts were injected exclusively into each individual group. The Avg@8 accuracy is reported in Table 10.

Injection	GSM8K	MATH500
GRPO (Baseline)	76.95	82.75
Bottom	77.15	83.10
Middle	77.58	83.65
Top	77.05	82.95
SCM (All Layers)	77.79	84.00

Table 10: Layer-wise SCM ablation results (Avg@8).

The results show that injecting soft concepts into middle layers provides competitive improvements, suggesting that these layers may offer an effective interface for incorporating continuous semantic information. Meanwhile, full-layer injection achieves the best overall performance, indicating that different layers can mutually benefit from continuous concepts during reasoning.