

000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 VARICoT: A VARIATIONAL FRAMWORK FOR IMPLICIT REASONING AS STRUCTURED LATENT INFERENCE

Anonymous authors

Paper under double-blind review

ABSTRACT

Chain-of-Thought (CoT) elicits remarkable capabilities in large language models but is fundamentally constrained by the low-bandwidth, sequential nature of text generation. Implicit CoT methods promise to accelerate by reasoning on latent space, yet they often rely on heuristic architectures and complex multi-stage training, lacking a unified, principled foundation. We introduce VARICOT, the first principled variational framework that formulates implicit reasoning as a structured probabilistic inference problem. VARICOT learns a continuous latent variable, Z , that represents the entire reasoning process, optimized via a single, unified evidence lower bound (ELBO) objective. Our key architectural innovation, Guided Latent Reasoning, treats Z as a global reasoning context that modulates the model's computations at every layer via cross-attention. This design decouples the abstract reasoning state from the linguistic realization, enabling high-bandwidth guidance without altering the standard autoregressive generation process. Implemented within a single Transformer and trained end-to-end with strategic control tokens, VARICOT offers flexible inference: either generating answers directly for a $>2.5x$ speedup or reproducing the full rationale when needed. On benchmarks like GSM8K and CommonsenseQA, VARICOT substantially improves upon or matches the accuracy of explicit CoT while drastically reducing latency, establishing a theoretically grounded and scalable paradigm for efficient reasoning.

1 INTRODUCTION

Large language models (LLMs) have demonstrated remarkable reasoning capabilities, particularly when guided by explicit chain-of-thought (CoT) prompting that verbalizes intermediate steps in natural language (Wei et al., 2022; Kojima et al., 2022). While effective, this paradigm imposes a fundamental bottleneck: models capable of manipulating thousands of dimensions in their internal hidden states are forced to reason through the narrow, discrete channel of token-by-token generation. (Zhu et al., 2025). This mismatch between the model's high-dimensional computational capacity and the low-bandwidth nature of text couples the act of "thinking" with the act of "writing," incurring significant computational overhead and hindering end-to-end optimization of the reasoning process itself.

The limitations of explicit CoT have motivated a shift toward *latent reasoning*—performing multi-step inference entirely within the model's continuous hidden space without generating intermediate tokens. Large language models (LLMs) can solve complex problems by externalizing their reasoning process as a "chain of thought" (CoT) in natural language. However, this paradigm introduces a fundamental paradox: models capable of manipulating thousands of dimensions in their internal hidden states are forced to reason through the narrow, discrete channel of token-by-token generation.

The limitations of explicit CoT have spurred the development of latent reasoning—performing multi-step inference within the model's continuous hidden space without generating intermediate tokens (Geiping et al., 2025; Ruan et al., 2025; Hao et al., 2024; Shen et al., 2025). While promising, current approaches represent a zoo of disparate, often heuristic solutions. They range from injecting "latent tokens" into the input sequence (vertical reasoning) to modifying hidden states across layers (horizontal reasoning) (Zhu et al., 2025). More critically, these methods frequently depend on complex, multi-stage training pipelines, such as distillation from an explicit CoT "teacher" model or reliance on external memory modules (Dehghani et al., 2018; Sun et al., 2024; Behrouz et al.,

054 2024; Shen et al., 2025). This architectural fragmentation and lack of a shared theoretical objective
 055 prevent true end-to-end optimization, making the learned reasoning processes opaque and difficult to
 056 generalize. What is missing is a unified framework that combines the efficiency of latent reasoning
 057 with a principled learning objective.

058 In this work, we introduce **VARICOT (Variational Implicit Chain-of-Thought)**, a unified frame-
 059 work that addresses these challenges by reformulating implicit reasoning as a problem of structured
 060 probabilistic inference. Instead of heuristics, we treat the unobserved reasoning process as a continu-
 061 ous, structured latent variable Z . We then derive a single evidence lower bound (ELBO) objective
 062 that jointly learns to: (1) infer Z from the problem context, (2) generate the final answer conditioned
 063 on Z , and (3) optionally reconstruct the explicit reasoning trace for interpretability. This principled,
 064 variational foundation allows for end-to-end training of the entire reasoning architecture within a
 065 single model.

066 VARICOT is realized through two synergistic innovations. First, we propose **Guided Latent**
 067 **Reasoning**, a novel architectural paradigm that cleanly decouples the latent reasoning state (Z) from
 068 the token-level representations. In our design, Z acts as an external, global context that is shared
 069 across all Transformer layers. At each layer, the model uses cross-attention to query this latent
 070 context, allowing the reasoning state to guide linguistic processing without being entangled in the
 071 residual stream. This synthesizes the representational power of continuous states with the structural
 072 integrity of autoregressive models. Second, we employ **Strategic Control Tokens** Goyal et al.
 073 (2024); Wang et al. (2024) (e.g., `<prior>`, `<posterior>`) to manage the different probabilistic
 074 operations (prior sampling, posterior inference, and conditional generation) within a single, standard
 075 autoregressive pass. This lightweight mechanism eliminates the need for multi-stage pipelines or
 076 architectural modifications, enabling seamless, scalable, and end-to-end training.

077 We evaluate variCoT across arithmetic, symbolic, and commonsense reasoning benchmarks. Our
 078 framework consistently outperforms strong explicit CoT and latent baselines, while exhibiting
 079 superior sample efficiency and robustness to prompt perturbations. Ablation studies confirm that the
 080 variational objective is essential: it not only improves performance but also encourages disentangled,
 081 interpretable latent representations that align with ground-truth reasoning steps.

082 In summary, our contributions are:

- 084 • **variCoT, A Unified Variational Framework:** We are the first to formalize implicit CoT
 085 reasoning within a principled variational inference framework, optimizing a joint evidence
 086 lower bound (ELBO) in an end-to-end fashion.
- 087 • **Guided Latent Reasoning:** We introduce a novel architecture that decouples reasoning
 088 and language by using the latent reasoning state as a cross-attentional query, providing
 089 high-bandwidth guidance to every layer of the Transformer.
- 090 • **Single-Model, End-to-End Training:** We demonstrate how strategic control tokens can
 091 embed the entire variational machinery into a standard Transformer, eliminating the need
 092 for complex, multi-stage training pipelines like distillation.
- 093 • **Improved Performance:** We will show through extensive experiments on arithmetic, com-
 094 monsense, and symbolic reasoning tasks that VARICOT achieves state-of-the-art accuracy
 095 and sample efficiency while offering significant inference speedups ($>2.5x$) over explicit
 096 CoT methods.

100 2 METHODOLOGY

101 We introduce variCoT, a unified variational framework for implicit Chain-of-Thought reasoning
 102 that addresses fundamental limitations in existing latent reasoning approaches. While methods
 103 like explicit CoT are constrained by discrete token sequences and latent approaches often rely on
 104 heuristic architectures or multi-stage training, variCoT provides a principled probabilistic foundation
 105 for learning continuous reasoning traces within a single Transformer. This section formalizes our
 106 approach through a structured generative model, derives its training objective via variational inference,
 107 and demonstrates how each component overcomes key challenges in latent reasoning.

108 2.1 BACKGROUND AND NOTATION
109

110 We begin by establishing the formal setting for reasoning in large language models. Let $X^q =$
111 (x_1^q, \dots, x_n^q) denote the input question token sequence, $Y^r = (y_1^r, \dots, y_m^r)$ the explicit reasoning chain,
112 and $Y^a = (y_1^a, \dots, y_k^a)$ the final answer. Standard autoregressive language models generate these
113 components sequentially using the factorization $p(Y^r, Y^a | X^q) = p(Y^r | X^q) \cdot p(Y^a | X^q, Y^r)$.

114 The fundamental limitation of this approach lies in the information bottleneck of discrete tokens.
115 Each token carries approximately 15 bits of information, while a single hidden state in modern LLMs
116 (e.g., 4096-dimensional) can encode 40,960 bits—a 2,700× increase in expressive capacity Zhu et al.
117 (2025). This observation has motivated latent reasoning methods that operate in continuous hidden
118 spaces. However, existing approaches such as Coconut Hao et al. (2024) and CODI Shen et al. (2025)
119 rely on deterministic recurrence or distillation pipelines, lacking proper uncertainty quantification
120 and end-to-end optimization.

121 variCoT addresses these limitations by introducing a sequence of continuous latent variables
122 $Z = (z_1, \dots, z_L)$ that serves as a compressed, stochastic representation of the reasoning process.
123 Unlike prior work, our framework formalizes Z within a generative model, enabling principled
124 variational inference and uncertainty-aware reasoning while maintaining full compatibility with
125 standard Transformer architectures.

126 2.2 THE VARICoT FRAMEWORK
127

128 variCoT is grounded in two key insights from the latent reasoning literature: (1) the expressive
129 advantage of continuous hidden states over discrete tokens, and (2) the functional specialization of
130 Transformer layers—shallow layers for representation, intermediate for transformation, and deep
131 for integration Skean et al. (2024); Gromov et al. (2024); Shi et al. (2024); Zhang et al. (2024). We
132 mirror this structure by letting Z encapsulate the full reasoning trajectory before branching into
133 separate decoders for reasoning and answer generation. We begin by establishing a general theoretical
134 foundation for variational reasoning without imposing any assumption:

135 **Theorem 2.1** (Evidence Lower Bound for Latent Reasoning). *For any joint distribution $p(Y^r, Y^a, Z |$
136 $X^q)$ and variational approximation $q_\phi(Z | X^q, Y^r, Y^a)$, the log marginal likelihood admits the
137 decomposition:*

$$138 \log p(Y^r, Y^a | X^q) = \mathcal{L}_{ELBO} + D_{KL}(q_\phi(Z | X^q, Y^r, Y^a) \| p(Z | X^q, Y^r, Y^a)),$$

139 where

$$140 \mathcal{L}_{ELBO} = \mathbb{E}_{q_\phi} \left[\log \frac{p(Y^r, Y^a, Z | X^q)}{q_\phi(Z | X^q, Y^r, Y^a)} \right].$$

141 *Proof.* The derivation follows from a variational decomposition of the log marginal likelihood,
142 leveraging the non-negativity of the Kullback-Leibler divergence. A complete derivation is provided
143 in Appendix A.1. \square

144 While Theorem 2.1 provides a general variational foundation, it presents two practical challenges
145 for reasoning applications. First, the KL divergence term requires access to the true posterior
146 $p(Z | X^q, Y^r, Y^a)$, which is intractable. Second, even with a variational approximation q_ϕ , the
147 posterior remains conditioned on both Y^r and Y^a , making it unusable during inference when
148 reasoning chains are unavailable.

149 To address these limitations, we introduce a structured generative model that enables tractable
150 optimization and practical deployment. Our approach is motivated by the observation that effective
151 reasoning requires a clean separation between abstract computation and linguistic realization.

152 **Assumption 2.2** (Latent Reasoning Mediation). *There exists a sequence of latent reasoning states
153 $Z = (z_1, \dots, z_L)$ such that, conditioned on the question X^q and Z , the explicit reasoning Y^r and
154 the answer Y^a are conditionally independent:*

$$155 Y^r \perp\!\!\!\perp Y^a | X^q, Z.$$

156 This assumption reflects the cognitive intuition that once the core reasoning process is complete, its
157 verbalization (Y^r) and final answer (Y^a) can be generated independently. It aligns with empirical

162 findings on layer-wise specialization in Transformers, where shallow layers handle surface features
 163 while deeper layers integrate semantic and inferential content.

164 Under Assumption 2.2, we obtain a tractable factorization of the joint distribution:

166 **Proposition 2.3** (variCoT Generative Factorization). *Under Assumption 2.2, the joint distribution*
 167 *over Y^r , Y^a , and Z given X^q factorizes as:*

$$168 \quad p_{\theta, \psi, \rho}(Y^r, Y^a, Z \mid X^q) = p_\psi(Y^r \mid X^q, Z) \cdot p_\rho(Y^a \mid X^q, Z) \cdot p_\theta(Z \mid X^q),$$

170 where $p_\theta(Z \mid X^q)$ is the prior over latent reasoning, and p_ψ , p_ρ model the generation of explicit
 171 reasoning and answer, respectively.

173 This factorization enables a computationally efficient training objective that bridges the theoretical
 174 ELBO with practical optimization:

175 **Theorem 2.4** (VariCoT Objective Decomposition). *Under the factorization in Proposition 2.3, the*
 176 *ELBO decomposes into three interpretable components:*

$$178 \quad \mathcal{L}_{ELBO} = \underbrace{\mathbb{E}_{q_\phi} [\log p_\psi(Y^r \mid X^q, Z)]}_{\mathcal{L}_{reasoning}} + \underbrace{\mathbb{E}_{q_\phi} [\log p_\rho(Y^a \mid X^q, Z)]}_{\mathcal{L}_{answer}} \\ 180 \\ 181 \quad - \underbrace{\beta \cdot D_{KL}(q_\phi(Z \mid X^q, Y^r, Y^a) \parallel p_\theta(Z \mid X^q))}_{\mathcal{L}_{KL}},$$

184 where $\beta > 0$ is a tunable regularization coefficient.

186 *Proof.* The decomposition follows from substituting the structured joint distribution into the ELBO
 187 and applying linearity of expectation. See Appendix A.2. \square

190 The decomposition in Theorem 2.4 provides a principled training objective where each term serves a
 191 distinct function. During training, the variational posterior $q_\phi(Z \mid X^q, Y^r, Y^a)$ absorbs all available
 192 information from both reasoning chains and answers. The KL regularization term \mathcal{L}_{KL} ensures that
 193 the prior $p_\theta(Z \mid X^q)$ learns to approximate this informed distribution, enabling effective inference
 194 when ground-truth reasoning chains are unavailable. This design allows the model to sample
 195 $Z \sim p_\theta(Z \mid X^q)$ at test time and generate Y^a directly, enabling efficient latent-only reasoning
 196 that bypasses explicit CoT generation while retaining the ability to reconstruct rationales when
 197 interpretability is required. The remaining terms provide complementary learning signals: $\mathcal{L}_{reasoning}$
 198 ensures the latent variable Z retains sufficient information to reconstruct explicit reasoning chains,
 199 serving as an interpretability anchor, while \mathcal{L}_{answer} drives task performance by ensuring Z encodes all
 200 necessary information for accurate final answers.

201 This formulation establishes variCoT as a probabilistically grounded framework for end-to-end train-
 202 able latent reasoning. Compared to heuristic or distillation-based approaches, our method provides
 203 theoretical guarantees through its ELBO foundation while addressing key limitations of prior work:
 204 it enables uncertainty-aware reasoning through distributional latent states, supports generalization via
 205 prior regularization, and maintains architectural flexibility through modular decoders.

207 3 IMPLEMENTING VARICoT: THE GUIDED LATENT TRANSFORMER

209 The variCoT framework proposes a unified variational objective for latent reasoning. To realize its
 210 full potential, we must address two practical challenges: (1) how to train all components—prior,
 211 posterior, reasoning decoder, and answer decoder—efficiently within a single model, and (2) how
 212 to represent and inject the latent variable Z to achieve high-bandwidth reasoning while maintaining
 213 architectural compatibility. We solve the first challenge through strategic control tokens that enable
 214 end-to-end training, and the second through guided latent reasoning, a novel architectural paradigm
 215 that synthesizes the strengths of existing approaches. The complete training and inference procedures
 are summarized in Algorithms 1 and 2 in the appendix.

216
217

3.1 STRATEGIC CONTROL TOKENS: END-TO-END SINGLE-MODEL TRAINING

218
219
220
221
222
223

A major limitation of existing latent reasoning frameworks is their reliance on multi-stage pipelines (Hao et al., 2024)—such as knowledge distillation (Shen et al., 2025), external encoders for discretization (Su et al., 2025), or persistent memory modules (Gao et al., 2024)—which fragment the computational graph, increase memory overhead, and hinder scalability within standard autoregressive architectures. We address this by introducing *strategic control tokens* that enable end-to-end variational inference within a single Transformer.

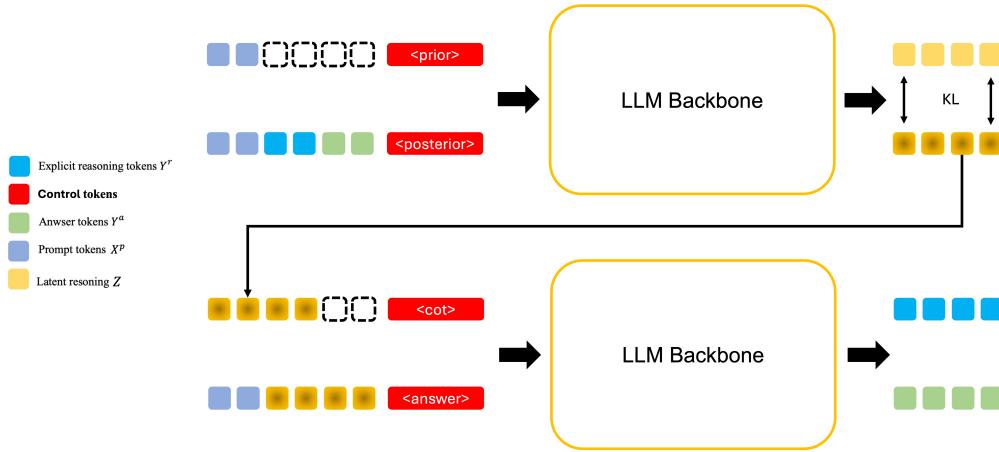
224
225
226
227
228
229
230
231
232
233
234
235
236
237
238

Figure 1: Training data flow in variCoT. Control tokens condition distinct probabilistic operations within a single forward pass. The latent variable Z is sampled from the posterior and used to guide decoding. All components share parameters, enabling end-to-end training. Our approach builds on training-induced recurrence (Goyal et al., 2024; Wang et al., 2024), where structured token sequences induce specialized computational roles without architectural modification. We extend this idea to variational learning by embedding the full generative and inference machinery into a unified sequence via functionally specialized tokens.

245
246
247
248
249
250
251

During training (Figure 1), a single forward pass processes the input question X_q , ground-truth reasoning trace Y_r , answer Y_a , and control tokens $\langle \text{prior} \rangle$ and $\langle \text{posterior} \rangle$. The hidden state at $\langle \text{posterior} \rangle$ parameterizes the approximate posterior $q_\phi(Z | X_q, Y_r, Y_a)$, from which Z is sampled and routed to the reasoning and answer decoders. Crucially, all components share the same Transformer parameters, enabling uninterrupted gradient flow. At inference, the model samples Z from the prior $p_\theta(Z | X_q)$ and generates outputs autoregressively. For complete implementation details including token specifications and training protocols, see Appendix A.3.

252
253
254
255

By unifying probabilistic operations through token-level control, our method achieves full compatibility with pretrained LLMs while supporting expressive, uncertainty-aware reasoning—resolving key scalability and modularity challenges identified in recent latent reasoning literature (Sui et al., 2025).

256
257

3.2 LATENT REPRESENTATION PARADIGMS: VERTICAL, HORIZONTAL, AND HYBRID APPROACHES

258
259
260
261
262
263
264
265

Following the taxonomy of latent reasoning frameworks Zhu et al. (2025), we formalize three paradigms for representing the latent variable Z in variational reasoning. Each defines a distinct architectural pathway for coupling latent reasoning states with autoregressive language modeling. **Vertical Paradigm: Discrete Latent Tokens** In the vertical paradigm, the latent variable is instantiated as a sequence of discrete tokens $Z = (z_1, \dots, z_S)$, where each z_s is drawn from a learned categorical distribution over a fixed latent vocabulary \mathcal{V} . These tokens are embedded and concatenated with the input token embeddings to form a joint sequence processed autoregressively. The architecture is defined by:

$$z_s \sim \text{Categorical}(\pi_\theta(x_{\leq t}, z_{<s})) \quad \forall s \in \{1, \dots, S\}, \quad (1)$$

$$\mathbf{H}_{\text{input}} = \text{Concat}(e(x_1), \dots, e(x_T), e(z_1), \dots, e(z_S)), \quad (2)$$

266
267
268
269

where $e(\cdot)$ denotes the token embedding function and π_θ is a parameterized policy conditioned on prior inputs and latent tokens. This formulation enables direct interpretability and intervention at

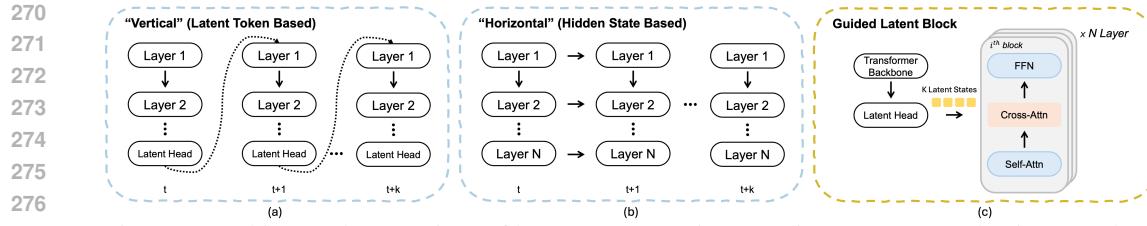


Figure 2: Architectural comparison of latent representation paradigms. (a) *Vertical*: Discrete tokens processed autoregressively. (b) *Horizontal*: Continuous hidden states injected into residual stream. (c) *Hybrid (Ours)*: Continuous latent states as per-layer guidance via cross-attention.

the token level. However, the information capacity of Z is limited by the discrete tokens, typically ~ 15 bits per token (Zhu et al., 2025), restricting the complexity of representable reasoning states and undermining the expressive potential of continuous latent spaces.

Horizontal Paradigm: Continuous Hidden States The horizontal paradigm identifies Z with a subset of the model’s internal continuous hidden states. Specifically, Z is extracted from the transformer layer and re-injected into subsequent layers. The architecture is formalized as:

$$Z = h_t^{(l)} \in \mathbb{R}^d, \quad (3)$$

$$\mathbf{H}_{\text{input}}^{(l+1)} = \text{Concat}(\mathbf{H}^{(l)}, Z), \quad (4)$$

where $h_t^{(l)}$ is the hidden state at layer l and position t , and $\mathbf{H}^{(l)} \in \mathbb{R}^{T \times d}$ denotes the full sequence of activations at that layer. This approach preserves high information bandwidth—each d -dimensional vector encodes $O(d)$ bits—but entangles reasoning states with linguistic representations. As a result, the model struggles to disentangle task-agnostic reasoning dynamics from surface-level language features, complicating regularization, interpretation, and cross-task generalization.

Hybrid Paradigm: Guided Latent Reasoning We propose **guided latent reasoning** (Figure 2 (c)), a novel architectural paradigm that addresses the limitations of both vertical and horizontal approaches by decoupling the latent reasoning state from the autoregressive token stream while enabling fine-grained, layer-specific influence. This design preserves the structural clarity of discrete tokens while leveraging the representational capacity of continuous hidden states.

The key innovation treats the latent variable $Z = \{Z_1, \dots, Z_K\}$ as an external guidance bank $\mathbf{Z} \in \mathbb{R}^{K \times d}$ that provides global contextual guidance. Inspired by conditioning mechanisms in Diffusion Transformers Peebles & Xie (2023), \mathbf{Z} is sampled once during training or inference and shared across all transformer layers:

$$\mathbf{Z} = \text{MLP}_{\text{latent}}([\mathbf{H}^{\text{backbone}}]) \in \mathbb{R}^{K \times d},$$

where $\mathbf{H}^{\text{backbone}}$ is obtained from the backbone transformer processing the input context.

Rather than interleaving Z with tokens or overwriting activations, we augment each transformer block with cross-attention where \mathbf{Z} serves as query and the self-attended representations provide keys and values:

$$\mathbf{H}_{\text{self}}^{(l)} = \text{SelfAttn} \left(\text{LayerNorm}(\mathbf{H}^{(l-1)}) \right) + \mathbf{H}^{(l-1)}, \quad (5)$$

$$\mathbf{H}_{\text{cross}}^{(l)} = \text{CrossAttn} \left(\text{LayerNorm}(\mathbf{Z}), \text{LayerNorm}(\mathbf{H}_{\text{self}}^{(l)}), \text{LayerNorm}(\mathbf{H}_{\text{self}}^{(l)}) \right), \quad (6)$$

$$\mathbf{H}_{\text{merged}}^{(l)} = \mathbf{H}_{\text{self}}^{(l)} + g_l \cdot \mathbf{H}_{\text{cross}}^{(l)}, \quad (7)$$

where g_l is a learnable gate that modulates guidance strength per layer.

This establishes a clean separation between the *reasoning trace* (evolving token representations) and *reasoning state* (external \mathbf{Z}). The adaptive gating g_l naturally aligns with transformer layer specialization—minimizing interference in shallow layers while amplifying reasoning influence in deeper layers Geva et al. (2020). Critically, since \mathbf{Z} resides outside the token sequence, it preserves full autoregressive compatibility without consuming sequence length or disrupting causal masking. This hybrid approach achieves an optimal balance: maintaining the expressive power of continuous latent spaces while providing precise architectural control over reasoning dynamics.

324

4 EXPERIMENTS

327 Table 1: Main results on mathematical and commonsense reasoning benchmarks. We compare our
328 **variCoT** variants against strong baselines across two model families. The best score for each dataset
329 is in **bold**. The best score among our proposed variants is underlined.

330 Model	331 GSM8k	331 GSM8k-NL	331 CommonsenseQA	331 SVAMP	331 GSM-Hard	331 MultiA
GPT-2						
333 CoT-SFT	44.1	34.8	36.9	41.8	9.8	90.7
334 No-CoT-SFT	19.1	19.1	20.5	16.4	4.3	41.1
335 Pause-CoT-SFT	16.4	16.4	-	14.8	4.1	39.2
336 iCoT	30.1	3.2	26.2	29.4	5.7	55.5
337 Coconut	34.1	24.9	38.6	36.4	7.9	82.2
338 CODI	43.7	35.3	44.0	42.9	9.9	92.8
339 variCoT-Vertical	38.5	31.5	<u>38.0</u>	37.0	9.2	84.8
340 variCoT-Horizontal	39.6	32.0	<u>37.3</u>	37.8	<u>9.4</u>	83.6
341 variCoT-Guided	<u>43.9</u>	<u>35.4</u>	37.9	<u>42.6</u>	<u>9.4</u>	<u>91.5</u>
LLaMA3.2-1b						
343 CoT-SFT	61.6	54.1	68.2	66.7	15.8	99.3
344 No-CoT-SFT	30.9	30.9	74.9	44.1	7.1	70.9
345 Pause-CoT-SFT	28.1	28.1	-	41.2	6.7	65.3
346 iCoT	19.0	15.2	72.6	40.9	4.4	39.0
347 Coconut	45.3	27.2	60.6	48.8	9.9	90.1
348 CODI	55.6	49.7	74.0	61.1	12.8	96.1
349 variCoT-Vertical	51.3	42.8	77.2	61.3	13.3	94.1
350 variCoT-Horizontal	51.5	43.0	76.4	60.8	13.1	94.3
351 variCoT-Guided	<u>57.5</u>	<u>53.75</u>	<u>78.1</u>	<u>65.2</u>	<u>15.6</u>	<u>98.5</u>

352 We conducted experiments on both GPT2 Radford et al. (2019) and LLaMA3.2-1b Grattafiori et al.
353 (2024) to validate the generalizability of our method across different foundation models. For training,
354 we employed the AdamW (Loshchilov & Hutter, 2017) optimizer with a learning rate of 5×10^{-5} ,
355 incorporating 10% warm-up steps followed by linear decay. The GPT2 model (Radford et al., 2019)
356 was trained for 30 epochs, while LLaMA3.2-1b (Grattafiori et al., 2024) was trained for 15 epochs,
357 both with an effective batch size of 256. Regarding hyperparameter configuration, we selected 6
358 latent reasoning embeddings with $\beta = 0.01$ to align with other methods in the baseline; further
359 hyperparameter analysis can be found in our ablation studies. To ensure reproducibility, we set a fixed
360 random seed (seed=42) for all experiments, and each reported result represents a single run under
361 this controlled setting. All experiments were performed on an `m1.p5en.48xlarge` instance of
362 Amazon Elastic Compute Cloud, which includes 8 NVIDIA H200 (141GB) GPUs, using PyTorch
363 2.6 (Paszke et al., 2019) as the deep learning framework.

364 **Dataset** Following Shen et al. (2025), we evaluate **variCoT** on six public datasets, categorized
365 into in-domain and out-of-domain (OOD) settings for evaluation. We use three datasets for in-
366 domain evaluation. **GSM8k-Aug** (Deng et al., 2023) is a math reasoning dataset of 385K samples,
367 augmented from GSM8K (Cobbe et al., 2021) using GPT-4, with structured mathematical expressions
368 as rationales. **GSM8k-Aug-NL** Shen et al. (2025) is a variant of GSM8k-Aug where the reasoning
369 process is presented in natural language. **CommonsenseQA-CoT** (Shen et al., 2025), which extends
370 the original CommonsenseQA (Talmor et al., 2018) with Chain-of-Thought (CoT) annotations that
371 were generated using GPT-4o-mini and filtered for correctness. To evaluate robustness, we train on
372 GSM8k-Aug and test on three OOD datasets. **SVAMP** (Patel et al., 2021) is an elementary school
373 math word problem dataset. **GSM-HARD** (Gao et al., 2023) is a more challenging version of the
374 GSM8K test set with an expanded value range. **MultiArith** (Roy & Roth, 2015) is a multi-step
375 arithmetic word problem dataset from MAWPS (Koncel-Kedziorski et al., 2016).

376 **Baselines** We compare our method, **variCoT**, against several strong baselines that explore explicit
377 and implicit reasoning: **CoT-SFT**, standard supervised fine-tuning (SFT) on explicit chain-of-thought

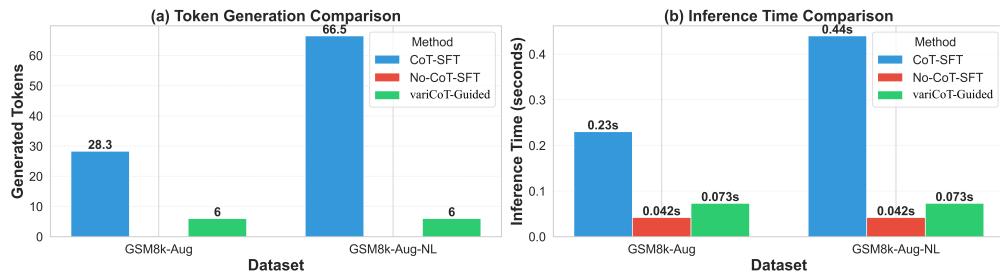
378 demonstrations, where the model generates the reasoning process before the final answer at inference;
 379 **No-CoT-SFT**, standard SFT on question-answer pairs only, without explicit reasoning steps; **Pause-**
 380 **CoT-SFT** (Goyal et al., 2024), SFT with special `<pause>` tokens inserted before the answer to
 381 encourage implicit reasoning (we use 6 for a fair comparison); **iCoT** (Deng et al., 2024), a strategy
 382 that internalizes reasoning by gradually removing the explicit CoT during training to ultimately output
 383 only the final answer; and **COCONUT** (Hao et al., 2024), a method that also internalizes the CoT,
 384 but replaces it with learned implicit reasoning tokens instead of deleting it. **CODI** (Shen et al., 2025),
 385 a method that also internalizes the CoT, but uses a distillation framework to compress the knowledge
 386 from an explicit CoT (teacher) process into a series of continuous thought tokens (student).

388 4.1 MAIN RESULTS

390 Table 1 presents comprehensive evaluations across mathematical and commonsense reasoning bench-
 391 marks. Our variCoT framework demonstrates strong performance across both GPT-2 and LLaMA3.2-
 392 1B model families, consistently matching or exceeding the accuracy of explicit CoT-SFT while
 393 offering significant efficiency gains.

394 On GPT-2, variCoT achieves 43.9% accuracy on GSM8K and 91.5% on MultiArith, performing com-
 395 petitively with explicit CoT-SFT (44.1% and 90.7% respectively) while significantly outperforming
 396 other implicit reasoning methods. The framework shows particular strength on out-of-domain gener-
 397 alization, achieving 42.6% on SVAMP and 9.4% on GSM-HARD, demonstrating robust reasoning
 398 capabilities without explicit intermediate token generation.

399 The performance advantage scales effectively to the larger LLaMA3.2-1B model, where variCoT
 400 achieves 57.5% on GSM8K and 98.5% on MultiArith—closely approaching CoT-SFT performance
 401 (61.6% and 99.3%) while offering the efficiency benefits of latent reasoning. Notably, our method
 402 shows superior commonsense reasoning capabilities, achieving 78.1% on CommonsenseQA-CoT,
 403 outperforming all baselines including explicit CoT-SFT (68.2%).



413 Figure 3: Inference efficiency of different methods in different datasets. The left side shows the
 414 average number of CoTs generated during the inference process, and the right side shows the average
 415 duration of complete inference, with GPT-2 Small as the base model.

416 In terms of inference efficiency, variCoT demonstrates significant advantages. As shown in Figure 3,
 417 our method reduces token generation by approximately 80-90% compared to CoT-SFT, requiring
 418 only 6 latent tokens instead of lengthy reasoning chains. This translates to a 70-80% reduction in
 419 inference time (0.073s vs. 0.32s for CoT-SFT on GSM8K), while maintaining competitive accuracy.
 420 Although slightly slower than No-CoT-SFT, this small efficiency sacrifice is exchanged for substantial
 421 performance gains, providing an excellent balance between efficiency and reasoning capability.

423 Table 2: CoT Reconstruction Quality Evaluation

425 Model	426 GSM8K-Aug		426 GSM8K-NL-Aug	
	427 ROUGE-1	427 BLEU-1	427 ROUGE-1	427 BLEU-1
GPT-2	0.69	0.66	0.63	0.62
LLaMA-1B	0.78	0.72	0.72	0.69

430 A key advantage of variCoT is its reversible reasoning capability. As shown in Table 2, our model
 431 achieves high reconstruction fidelity with ROUGE-1 scores of 0.69 (GPT-2) and 0.78 (LLaMA-1B)
 432 on GSM8K, indicating that the latent embeddings effectively capture essential reasoning information.

432 This provides significant interpretability advantages over other implicit CoT methods, as demonstrated
 433 by the reconstruction examples in Figure 7.
 434

435 **4.2 ABLATION STUDIES**
 436

437 We conduct systematic ablations to understand the impact of key architectural choices and hyper-
 438 parameters. First, we compare the three latent representation paradigms introduced in Section 3.2.
 439 The guided latent reasoning approach consistently outperforms both vertical (discrete token) and
 440 horizontal (continuous hidden state) variants across all benchmarks. On GPT-2, the guided paradigm
 441 achieves 43.9% on GSM8K compared to 38.5% for vertical and 39.6% for horizontal approaches.
 442 This advantage is even more pronounced on LLaMA3.2-1B, where the guided approach reaches
 443 57.5% versus 51.3% and 51.5% for the alternatives. The results validate our architectural design that
 444 decouples reasoning guidance from linguistic processing through cross-attention mechanisms.
 445

446 Figure 6 shows the sensitivity analysis of the number of latent reasoning embeddings. We find optimal
 447 performance scales with number of latent reasoning embeddings, consistent with findings in prior
 448 work Zhu et al. (2025). This configuration balances representational capacity with training stability,
 449 providing sufficient bandwidth for complex reasoning while avoiding overfitting. Performance
 450 degrades with fewer embeddings due to limited expressivity, while excessive embeddings introduce
 451 noise and optimization challenges.
 452

453 **5 RELATED WORK**
 454

455 Chain-of-Thought (CoT) prompting has established a powerful paradigm for multi-step reasoning in
 456 large language models by verbalizing intermediate steps (Wei et al., 2022). However, constraining
 457 reasoning to a discrete token space introduces significant latency and limits expressive power,
 458 motivating a shift towards *implicit* or *latent* CoT, where reasoning occurs in the model’s continuous
 459 hidden states (Zhu et al., 2025). Current latent reasoning methods primarily fall into two categories:
 460 *vertical* approaches that increase effective model depth by iteratively refining activations within a fixed
 461 set of layers (Geiping et al., 2025; Mohtashami et al., 2023), and *horizontal* approaches that expand
 462 temporal context by propagating compressed hidden states over time (Dao & Gu, 2024; Behrouz
 463 et al., 2024). While these methods enhance reasoning, they often require specialized architectures or
 464 entangle reasoning states with linguistic representations.
 465

466 A parallel line of work induces latent reasoning capabilities through specialized training objectives on
 467 standard Transformer architectures. These strategies include using special pause tokens to encourage
 468 implicit computation (Goyal et al., 2024), progressively internalizing explicit CoT steps during
 469 fine-tuning (Deng et al., 2024), or compressing natural language rationales into continuous thought
 470 vectors via knowledge distillation (Hao et al., 2024; Shen et al., 2025). Although effective, these
 471 training-induced methods often depend on multi-stage pipelines or heuristic objectives, lacking a
 472 unified, end-to-end optimization framework.
 473

474 Variational inference, while a cornerstone of generative modeling, remains nascent in the context of
 475 continuous latent reasoning. Prior works have either relied on discrete latent tokens (Su et al., 2025)
 476 or learned compressed reasoning traces without a principled probabilistic foundation (Zhang et al.,
 477 2025), failing to provide a robust framework for structured stochastic inference.
 478

479 Our work variCoT, addresses these gaps by proposing a unified variational framework that formalizes
 480 latent reasoning as principled stochastic inference. Unlike prior methods, variCoT is optimized
 481 end-to-end via a single, theoretically grounded evidence lower bound (ELBO) objective within a
 482 standard Transformer. It introduces a *guided latent reasoning* mechanism that synthesizes the benefits
 483 of vertical depth and horizontal recurrence, using cross-attention to decouple abstract reasoning from
 484 its linguistic realization. This unique design enables efficient, latent-only inference for fast decoding
 485 while preserving the ability to generate explicit CoT for interpretability—a critical capability not
 486 offered by previous implicit reasoning methods.
 487

488 **6 CONCLUSION**
 489

490 Chain-of-Thought reasoning improves LLM performance but incurs significant computational over-
 491 head through sequential token generation. Existing implicit CoT methods rely on heuristic archi-
 492

486 tectures and multi-stage training, lacking principled optimization. We introduce variCoT, a unified
 487 variational framework that overcomes these limitations through an evidence lower bound objective,
 488 formalizing latent reasoning traces as continuous stochastic variables.
 489

490 Our framework combines strategic control tokens for end-to-end training with guided latent reasoning
 491 that decouples abstract computation from linguistic realization. Experiments show variCoT matches
 492 or exceeds explicit CoT accuracy while providing $2.5 \times$ faster inference and reversible reasoning
 493 capability. This establishes a theoretically grounded, scalable approach to efficient reasoning that
 494 bridges continuous latent spaces with autoregressive generation.
 495

REFERENCES

497 Ali Behrouz, Peilin Zhong, and Vahab Mirrokni. Titans: Learning to memorize at test time. *arXiv preprint*
 498 *arXiv:2501.00663*, 2024.

499 Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert,
 500 Jerry Tworek, Jacob Hilton, Reiichiro Nakano, et al. Training verifiers to solve math word problems. *arXiv*
 501 *preprint arXiv:2110.14168*, 2021.

502 Tri Dao and Albert Gu. Transformers are SSMs: Generalized models and efficient algorithms through structured
 503 state space duality. *arXiv preprint arXiv:2405.21060*, 2024.

504 Mostafa Dehghani, Stephan Gouws, Oriol Vinyals, Jakob Uszkoreit, and Łukasz Kaiser. Universal transformers.
 505 *arXiv preprint arXiv:1807.03819*, 2018.

506 Yuntian Deng, Kiran Prasad, Roland Fernandez, Paul Smolensky, Vishrav Chaudhary, and Stuart Shieber.
 507 Implicit chain of thought reasoning via knowledge distillation. *arXiv preprint arXiv:2311.01460*, 2023.

509 Yuntian Deng, Yejin Choi, and Stuart Shieber. From explicit CoT to implicit CoT: Learning to internalize CoT
 510 step by step. *arXiv preprint arXiv:2405.14838*, 2024.

511 Luyu Gao, Aman Madaan, Shuyan Zhou, Uri Alon, Pengfei Liu, Yiming Yang, Jamie Callan, and Graham
 512 Neubig. Pal: Program-aided language models. In *International Conference on Machine Learning*, pp.
 513 10764–10799. PMLR, 2023.

514 Yihang Gao, Chuanyang Zheng, Enze Xie, Han Shi, Tianyang Hu, Yu Li, Michael K. Ng, Zhenguo Li, and
 515 Zhaoqiang Liu. Algoformer: An efficient transformer framework with algorithmic structures. *arXiv preprint*
 516 *arXiv:2402.13572*, 2024.

517 Jonas Geiping, Sean McLeish, Neel Jain, John Kirchenbauer, Siddharth Singh, Brian R. Bartoldson, Bhavya
 518 Kailkhura, Abhinav Bhatele, and Tom Goldstein. Scaling up test-time compute with latent reasoning: A
 519 recurrent depth approach. *arXiv preprint arXiv:2502.05171*, 2025.

520 Mor Geva, Roei Schuster, Jonathan Berant, and Omer Levy. Transformer feed-forward layers are key-value
 521 memories. *arXiv preprint arXiv:2012.14913*, 2020.

523 Sachin Goyal, Ziwei Ji, Ankit Singh Rawat, Aditya Krishna Menon, Sanjiv Kumar, and Vaishnav Nagarajan.
 524 Think before you speak: Training language models with pause tokens. In *The Twelfth International Conference*
 525 *on Learning Representations*, 2024. URL <https://openreview.net/forum?id=ph04CRkPdC>.

526 Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle,
 527 Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, et al. The llama 3 herd of models. *arXiv*
 528 *preprint arXiv:2407.21783*, 2024.

529 Andrey Gromov, Kushal Tirumala, Hassan Shapourian, Paolo Glorioso, and Daniel A. Roberts. The unreasonable
 530 ineffectiveness of the deeper layers. *arXiv preprint arXiv:2403.17887*, 2024. URL <https://arxiv.org/abs/2403.17887>.

532 Shibo Hao, Sainbayar Sukhbaatar, DiJia Su, Xian Li, Zhiting Hu, Jason Weston, and Yuandong Tian. Training
 533 large language models to reason in a continuous latent space. *arXiv preprint arXiv:2412.06769*, 2024.

534 Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. Large language
 535 models are zero-shot reasoners. *Advances in neural information processing systems*, 35:22199–22213, 2022.

536 Rik Koncel-Kedziorski, Subhro Roy, Aida Amini, Nate Kushman, and Hannaneh Hajishirzi. MAWPS: A math
 537 word problem repository. In Kevin Knight, Ani Nenkova, and Owen Rambow (eds.), *Proceedings of the*
 538 *2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human*
 539 *Language Technologies*, pp. 1152–1157, San Diego, California, June 2016. Association for Computational
 Linguistics. doi: 10.18653/v1/N16-1136. URL <https://aclanthology.org/N16-1136/>.

540 Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*,
 541 2017.

542 Amirkeivan Mohtashami, Matteo Pagliardini, and Martin Jaggi. CoTformer: A chain-of-thought driven
 543 architecture with budget-adaptive computation cost at inference. *arXiv preprint arXiv:2310.10845*, 2023.

544 Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen,
 545 Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep
 546 learning library. *Advances in neural information processing systems*, 32, 2019.

547 Arkil Patel, Satwik Bhattacharya, and Navin Goyal. Are NLP models really able to solve simple math word
 548 problems? In Kristina Toutanova, Anna Rumshisky, Luke Zettlemoyer, Dilek Hakkani-Tur, Iz Beltagy,
 549 Steven Bethard, Ryan Cotterell, Tanmoy Chakraborty, and Yichao Zhou (eds.), *Proceedings of the 2021*
 550 *Conference of the North American Chapter of the Association for Computational Linguistics: Human*
 551 *Language Technologies*, pp. 2080–2094, Online, June 2021. Association for Computational Linguistics.
 552 doi: 10.18653/v1/2021.naacl-main.168. URL <https://aclanthology.org/2021.naacl-main.168/>.

553 William Peebles and Saining Xie. Scalable diffusion models with transformers. In *Proceedings of the IEEE/CVF*
 554 *international conference on computer vision*, pp. 4195–4205, 2023.

555 Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language models are
 556 unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019.

557 Subhro Roy and Dan Roth. Solving general arithmetic word problems. In Lluís Márquez, Chris Callison-
 558 Burch, and Jian Su (eds.), *Proceedings of the 2015 Conference on Empirical Methods in Natural Language*
 559 *Processing*, pp. 1743–1752, Lisbon, Portugal, September 2015. Association for Computational Linguistics.
 560 doi: 10.18653/v1/D15-1202. URL <https://aclanthology.org/D15-1202/>.

561 Yangjun Ruan, Neil Band, Chris J Maddison, and Tatsunori Hashimoto. Reasoning to learn from latent thoughts.
 562 *arXiv preprint arXiv:2503.18866*, 2025.

563 Zhenyi Shen, Hanqi Yan, Linhai Zhang, Zhanghao Hu, Yali Du, and Yulan He. Codi: Compressing chain-of-
 564 thought into continuous space via self-distillation. *arXiv preprint arXiv:2502.21074*, 2025.

565 Guangyuan Shi, Zexin Lu, Xiaoyu Dong, Wenlong Zhang, Xuanyu Zhang, Yujie Feng, and Xiao-Ming Wu.
 566 Understanding layer significance in LLM alignment. *arXiv preprint arXiv:2410.17875*, 2024.

567 Oscar Skean, Md Rifat Arefin, Yann LeCun, and Ravid Shwartz-Ziv. Does representation matter? exploring
 568 intermediate layers in large language models. *arXiv preprint arXiv:2412.09563*, 2024.

569 DiJia Su, Hanlin Zhu, Yingchen Xu, Jiantao Jiao, Yuandong Tian, and Qinqing Zheng. Token assorted: Mixing
 570 latent and text tokens for improved language model reasoning. *arXiv preprint arXiv:2502.03275*, 2025.

571 Yang Sui, Yu-Neng Chuang, Guanchu Wang, Jiamu Zhang, Tianyi Zhang, Jiayi Yuan, Hongyi Liu, Andrew Wen,
 572 Hanjie Chen, Xia Hu, et al. Stop overthinking: A survey on efficient reasoning for large language models.
 573 *arXiv preprint arXiv:2503.16419*, 2025.

574 Yu Sun, Xinhao Li, Karan Dalal, Jiarui Xu, Arjun Vikram, Genghan Zhang, Yann Dubois, Xinlei Chen, Xiaolong
 575 Wang, Sanmi Koyejo, et al. Learning to (learn at test time): RNNs with expressive hidden states. *arXiv*
 576 *preprint arXiv:2407.04620*, 2024.

577 Alon Talmor, Jonathan Herzig, Nicholas Lourie, and Jonathan Berant. Commonsenseqa: A question answering
 578 challenge targeting commonsense knowledge. *arXiv preprint arXiv:1811.00937*, 2018.

579 Xinyi Wang, Lucas Caccia, Oleksiy Ostapenko, Xingdi Yuan, William Yang Wang, and Alessandro Sordoni.
 580 Guiding language model reasoning with planning tokens. In *First Conference on Language Modeling*, 2024.
 581 URL <https://openreview.net/forum?id=wi9IfffRhVM>.

582 Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V. Le, Denny Zhou, et al.
 583 Chain-of-thought prompting elicits reasoning in large language models. In *Advances in neural information*
 584 *processing systems*, volume 35, pp. 24824–24837, 2022.

585 Jintian Zhang, Yuqi Zhu, Mengshu Sun, Yujie Luo, Shuohei Qiao, Lun Du, Da Zheng, Huajun Chen, and Ningyu
 586 Zhang. Lightthinker: Thinking step-by-step compression. *arXiv preprint arXiv:2502.15589*, 2025.

587 Yang Zhang, Yanfei Dong, and Kenji Kawaguchi. Investigating layer importance in large language models.
 588 *arXiv preprint arXiv:2409.14381*, 2024.

589 Rui-Jie Zhu, Tianhao Peng, Tianhao Cheng, Xingwei Qu, Jinfang Huang, Dawei Zhu, Hao Wang, Kaiwen Xue,
 590 Xuanliang Zhang, Yong Shan, et al. A survey on latent reasoning. *arXiv preprint arXiv:2507.06203*, 2025.

594 **A APPENDIX**

596 **A.1 PROOF OF THEOREM 1: EVIDENCE LOWER BOUND FOR VARIOT**

598 *Proof.* We begin by deriving the ELBO for the marginal log-likelihood $\log p(Y^r, Y^a | X^q)$. Introducing the
599 variational posterior $q_\phi(Z | X^q, Y^r, Y^a)$, we have:

$$\begin{aligned}
 600 \log p(Y^r, Y^a | X^q) &= \log \int p(Y^r, Y^a, Z | X^q) dZ \\
 601 &= \log \int q_\phi(Z | X^q, Y^r, Y^a) \frac{p(Y^r, Y^a, Z | X^q)}{q_\phi(Z | X^q, Y^r, Y^a)} dZ \\
 602 &\geq \mathbb{E}_{q_\phi} \left[\log \frac{p(Y^r, Y^a, Z | X^q)}{q_\phi(Z | X^q, Y^r, Y^a)} \right] \quad (\text{Jensen's inequality}) \\
 603 &= \mathbb{E}_{q_\phi} [\log p(Y^r, Y^a, Z | X^q) - \log q_\phi(Z | X^q, Y^r, Y^a)].
 \end{aligned}$$

604 Let $\mathcal{L}_{\text{ELBO}} = \mathbb{E}_{q_\phi} \left[\log \frac{p(Y^r, Y^a, Z | X^q)}{q_\phi(Z | X^q, Y^r, Y^a)} \right]$. We can rewrite the marginal likelihood as:

$$\begin{aligned}
 605 \log p(Y^r, Y^a | X^q) &= \mathbb{E}_{q_\phi} [\log p(Y^r, Y^a | X^q)] \\
 606 &= \mathbb{E}_{q_\phi} \left[\log \frac{p(Y^r, Y^a, Z | X^q)}{p(Z | X^q, Y^r, Y^a)} \right] \\
 607 &= \mathbb{E}_{q_\phi} \left[\log \frac{p(Y^r, Y^a, Z | X^q)}{q_\phi(Z | X^q, Y^r, Y^a)} \cdot \frac{q_\phi(Z | X^q, Y^r, Y^a)}{p(Z | X^q, Y^r, Y^a)} \right] \\
 608 &= \mathcal{L}_{\text{ELBO}} + D_{\text{KL}}(q_\phi(Z | X^q, Y^r, Y^a) \| p(Z | X^q, Y^r, Y^a)).
 \end{aligned}$$

609 Since the KL divergence is non-negative, we have $\log p(Y^r, Y^a | X^q) \geq \mathcal{L}_{\text{ELBO}}$, with equality if and only if
610 $q_\phi(Z | X^q, Y^r, Y^a) = p(Z | X^q, Y^r, Y^a)$. \square

611 **A.2 PROOF OF THEOREM 2: VARIOT OBJECTIVE DECOMPOSITION**

612 *Proof.* Starting from the ELBO expression in Theorem 1:

$$\mathcal{L}_{\text{ELBO}} = \mathbb{E}_{q_\phi} \left[\log \frac{p(Y^r, Y^a, Z | X^q)}{q_\phi(Z | X^q, Y^r, Y^a)} \right],$$

613 we substitute the joint distribution from Proposition 2.3:

$$\begin{aligned}
 614 \mathcal{L}_{\text{ELBO}} &= \mathbb{E}_{q_\phi} \left[\log \frac{p_\psi(Y^r | X^q, Z) \cdot p_\rho(Y^a | X^q, Z) \cdot p_\theta(Z | X^q)}{q_\phi(Z | X^q, Y^r, Y^a)} \right] \\
 615 &= \mathbb{E}_{q_\phi} [\log p_\psi(Y^r | X^q, Z) + \log p_\rho(Y^a | X^q, Z) + \log p_\theta(Z | X^q) - \log q_\phi(Z | X^q, Y^r, Y^a)].
 \end{aligned}$$

616 By linearity of expectation:

$$\begin{aligned}
 617 \mathcal{L}_{\text{ELBO}} &= \mathbb{E}_{q_\phi} [\log p_\psi(Y^r | X^q, Z)] + \mathbb{E}_{q_\phi} [\log p_\rho(Y^a | X^q, Z)] \\
 618 &\quad + \mathbb{E}_{q_\phi} \left[\log \frac{p_\theta(Z | X^q)}{q_\phi(Z | X^q, Y^r, Y^a)} \right].
 \end{aligned}$$

619 The third term can be rewritten as a KL divergence:

$$\mathbb{E}_{q_\phi} \left[\log \frac{p_\theta(Z | X^q)}{q_\phi(Z | X^q, Y^r, Y^a)} \right] = -D_{\text{KL}}(q_\phi(Z | X^q, Y^r, Y^a) \| p_\theta(Z | X^q)).$$

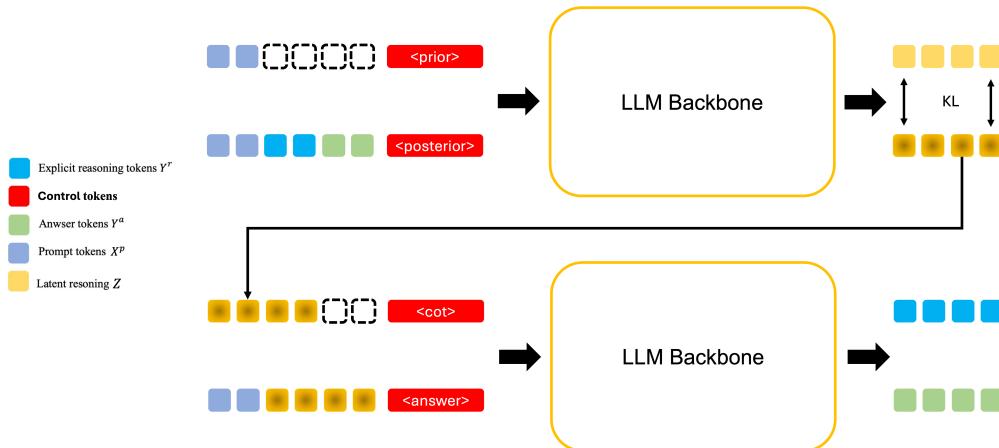
620 Thus, we obtain the final decomposition:

$$\begin{aligned}
 621 \mathcal{L}_{\text{ELBO}} &= \mathbb{E}_{q_\phi} [\log p_\psi(Y^r | X^q, Z)] + \mathbb{E}_{q_\phi} [\log p_\rho(Y^a | X^q, Z)] \\
 622 &\quad - D_{\text{KL}}(q_\phi(Z | X^q, Y^r, Y^a) \| p_\theta(Z | X^q)).
 \end{aligned}$$

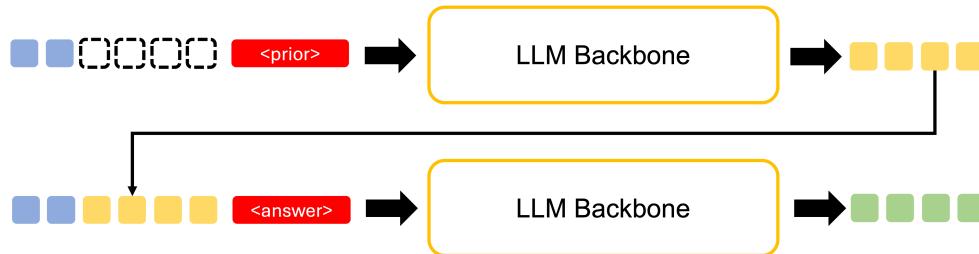
623 The introduction of the β coefficient follows the β -VAE framework to control the strength of the KL regularization
624 term, giving us the final objective:

$$\begin{aligned}
 625 \mathcal{L}_{\text{ELBO}} &= \underbrace{\mathbb{E}_{q_\phi} [\log p_\psi(Y^r | X^q, Z)]}_{\mathcal{L}_{\text{reasoning}}} + \underbrace{\mathbb{E}_{q_\phi} [\log p_\rho(Y^a | X^q, Z)]}_{\mathcal{L}_{\text{answer}}} \\
 626 &\quad - \underbrace{\beta \cdot D_{\text{KL}}(q_\phi(Z | X^q, Y^r, Y^a) \| p_\theta(Z | X^q))}_{\mathcal{L}_{\text{KL}}}.
 \end{aligned}$$

\square

648 A.3 STRATEGIC CONTROL TOKENS AND END-TO-END TRAINING PROTOCOL
649650 To ensure full reproducibility and clarity, we detail the complete token-level specification of variCoT’s control
651 mechanism, including exact input formats for all components and the inference procedure.652 **Control Token Specification** We define four functionally specialized control tokens that orchestrate
653 variational inference within a single Transformer pass. All components share the same set of parameters, and
654 gradients flow end-to-end through the entire sequence.
655656 • **Prior Network** $p_\theta(Z | X^q)$:657 Input sequence: <prompt> X^q </prompt> <prior>658 The latent variable $Z \in \mathbb{R}^{K \times d}$ is sampled from a distribution parameterized by the hidden state at the
<prior> token position. During inference, this is the sole source of Z .659 • **Posterior Network** $q_\phi(Z | X^q, Y^r, Y^a)$:660 Input sequence: <prompt> X^q </prompt> <cot> Y^r <eos> <answer> Y^a <eos>
<posterior>661 The approximate posterior is inferred from the hidden state at <posterior>, conditioned on both
662 the ground-truth reasoning chain Y^r and final answer Y^a . During training, Z is sampled from this
663 posterior.
664665 • **Reasoning Decoder** $p_\psi(Y^r | X^q, Z)$:666 Input sequence: <prompt> X^q </prompt> <latent> Z_1, \dots, Z_K </latent> <cot>667 Target sequence: Y^r followed by <eos>. The latent state Z is injected via the Guided Latent
668 Reasoning mechanism (Section 3.2). During training, Z is sampled from the posterior; during
669 inference, from the prior.
670671 • **Answer Decoder** $p_\phi(Y^a | X^q, Z)$:672 Input sequence: <prompt> X^q </prompt> <latent> Z_1, \dots, Z_K </latent> <answer>673 Target sequence: Y^a followed by <eos>.
674675 The <latent> token sequence serves as a placeholder that triggers the latent injection mechanism; its
676 embeddings are unused—the actual latent vectors Z are provided externally via cross-attention (Section 3.2).
677678 **Training Protocol** During training, we construct a single concatenated sequence:
679680 Figure 4: Training data flow in variCoT. Control tokens (<prior>, <posterior>, <cot>,
681 <answer>) condition distinct probabilistic operations within a single forward pass. The latent
682 variable Z is sampled from the posterior and routed to decoders via the <latent> token sequence.
683 Parameter sharing enables end-to-end gradient flow.
684685 $\text{<prompt>} X^q \text{ </prompt>} \text{<cot>} Y^r \text{ <eos>} \text{<answer>} Y^a \text{ <eos>} \text{<posterior>} \text{<latent>} Z_1, \dots, Z_K \text{ </latent>}$
686687 The model first processes the context up to <posterior> to infer $q_\phi(Z | X^q, Y^r, Y^a)$, samples Z , and then
688 uses this Z to condition the subsequent generation of Y^r and Y^a . The entire sequence is trained with standard
689 autoregressive language modeling loss, enabling end-to-end optimization.
690

702 **Inference Protocol** At inference time, no ground-truth reasoning trace or answer is available. The model
 703 proceeds in two stages: 1. **Latent sampling**: Process $\langle\text{prompt}\rangle X^q \langle/\text{prompt}\rangle \langle\text{prior}\rangle$ to obtain
 704



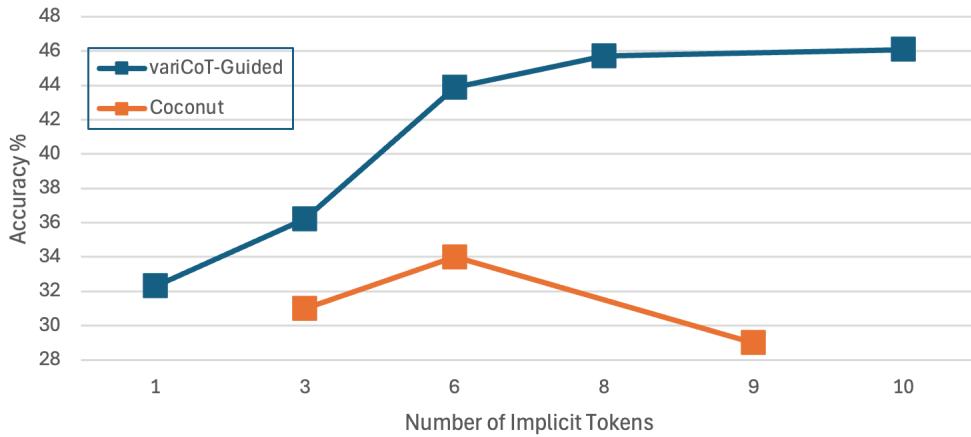
705
 706
 707
 708
 709
 710
 711
 712
 713
 714
 715
 716
 717
 718 Figure 5: Inference in variCoT. The prior network samples Z from $p_\theta(Z | X^q)$. The same latent
 719 state is then used to generate Y^r and Y^a sequentially. No ground-truth reasoning traces are required.
 720
 721

722 $p_\theta(Z | X^q)$, then sample Z . 2. **Autoregressive generation**: Using the sampled Z , generate either: - A
 723 reasoning chain: $\langle\text{prompt}\rangle X^q \langle/\text{prompt}\rangle \langle\text{latent}\rangle Z \langle/\text{latent}\rangle \langle\text{cot}\rangle \sim Y^r$ - A direct answer:
 724 $\langle\text{prompt}\rangle X^q \langle/\text{prompt}\rangle \langle\text{latent}\rangle Z \langle/\text{latent}\rangle \langle\text{answer}\rangle \sim Y^a$

725 Both generations are fully autoregressive and leverage the same latent state Z , enabling coherent, uncertainty-
 726 aware predictions.

727 This unified design eliminates the need for external encoders, distillation teachers, or non-autoregressive memory
 728 modules, while preserving full compatibility with standard transformer-based language models.

732 A.4 MORE EXPERIMENT RESULTS



752 Figure 6: The impact of different numbers of implicit reasoning embeddings and different lambda
 753 settings on performance under the GSM8K-Aug dataset, with GPT-2 Small as the base model.

756 A.5 EXPLICIT REASONING RECONSTRUCTION VISUALIZATION
757

758 Example 1

759 **Question:**760 Josh decides to try flipping a house. He buys a house for \$80,000 and then puts in \$50,000 in repairs. This
761 increased the value of the house by 150%. How much profit did he make?762 **Original Reasoning:**763 $<\!<\!80000+50000=130000\!>\!>\!<\!80000*1.5=120000\!>\!>\!<\!120000+80000=200000\!>\!>\!<\!200000-$
764 $130000=70000\!>\!>$ 765 **Reversible Latent Decoder Output:**766 $<\!<\!80000+50000=130000\!>\!>\!<\!150\%*80000=120000\!>\!>\!<\!80000+120000=200000\!>\!>\!<\!200000-$
767 $130000=70000\!>\!>$

768 Example 2

769 **Question:**770 Janet's ducks lay 16 eggs per day. She eats three for breakfast every morning and bakes muffins for her friends
771 every day with four. She sells the remainder at the farmers' market daily for \$2 per fresh duck egg. How much
772 in dollars does she make every day at the farmers' market?773 **Original Reasoning:**774 Janet sells $16 - 3 - 4 = 9$ duck eggs a day. She makes $9 * 2 = \$18$ every day at the farmer's market.775 **Reversible Latent Decoder Output:**776 Janet has 16 eggs daily, leaving $16 - 7 = 9$ eggs to sell. At \$2 per egg, she earns $9 \times \$2 = \18 .

777 Figure 7: Reconstructed performance of implicit reasoning.

778

779 How to interpret implicit reasoning has been a key challenge in this direction, especially for scenarios that
780 require explicit reasoning generation, such as mathematical proofs, logical attributions, etc. Although pure
781 implicit reasoning can achieve efficiency improvements, it cannot effectively obtain explicit reasoning processes.
782 Thanks to the *Reversible Latent Decoder*, our proposed variCoT can directly reconstruct explicit reasoning based
783 on implicit reasoning, which offers better interpretability compared to other implicit reasoning methods. As
784 shown in Fig. 7, implicit reasoning embeddings can be directly reconstructed into explicit reasoning by the
785 *Reversible Latent Decoder*, and can also support the generation of final answers in terms of thought paths, which
786 more intuitively demonstrates the superiority of our proposed method.787 A.6 TRAINING AND INFERENCE ALGORITHMS
788789 **Algorithm 1** variCoT Training Procedure

Require: Dataset $\mathcal{D} = \{(X^q, Y^r, Y^a)\}$, model parameters θ, ϕ, ψ, ρ
1: **while** not converged **do**
2: Sample batch $(X^q, Y^r, Y^a) \sim \mathcal{D}$
3: Construct input sequence: $S = [X^q, \text{<cot>}, Y^r, \text{<eos>}, \text{<answer>}, Y^a, \text{<eos>}, \text{<posterior>}]$
4: Compute hidden states: $\mathbf{H} = \text{Transformer}(S)$
5: Extract posterior parameters from $\mathbf{h}_{\text{<posterior>}}$: μ_ϕ, σ_ϕ
6: Sample latent: $Z \sim \mathcal{N}(\mu_\phi, \sigma_\phi^2)$
7: Construct decoding sequence: $S_{\text{dec}} = [X^q, \text{<latent>}, Z, \text{<cot>}]$
8: Compute reasoning loss: $\mathcal{L}_{\text{reasoning}} = -\log p_\psi(Y^r | X^q, Z)$
9: Construct answer sequence: $S_{\text{ans}} = [X^q, \text{<latent>}, Z, \text{<answer>}]$
10: Compute answer loss: $\mathcal{L}_{\text{answer}} = -\log p_\rho(Y^a | X^q, Z)$
11: Compute KL divergence: $\mathcal{L}_{\text{KL}} = D_{\text{KL}}(q_\phi(Z | X^q, Y^r, Y^a) || p_\theta(Z | X^q))$
12: Total loss: $\mathcal{L} = \mathcal{L}_{\text{reasoning}} + \mathcal{L}_{\text{answer}} + \beta \mathcal{L}_{\text{KL}}$
13: Update parameters via gradient descent: $\nabla_\theta \mathcal{L}$
14: **end while**

804

805

806

807

808

809

810
 811
 812
 813
 814
 815
 816
 817
 818
 819
 820
 821
 822
 823
 824
 825
 826
 827
 828
 829
 830
 831

Algorithm 2 variCoT Inference Procedure

Require: Input question X^q , trained model parameters θ, ρ
 834 1: Construct prior sequence: $S_{\text{prior}} = [X^q, \text{<prior>}]$
 835 2: Compute hidden states: $\mathbf{H} = \text{Transformer}(S_{\text{prior}})$
 836 3: Extract prior parameters from $\mathbf{h}_{\text{<prior>}}$: $\mu_\theta, \sigma_\theta$
 837 4: Sample latent: $Z \sim \mathcal{N}(\mu_\theta, \sigma_\theta^2)$
 838 5: Construct answer sequence: $S_{\text{ans}} = [X^q, \text{<latent>}, Z, \text{<answer>}]$
 839 6: Generate answer autoregressively: $Y^a \sim p_\rho(\cdot | X^q, Z)$
 840 7: **Optional:** Generate reasoning chain: $Y^r \sim p_\psi(\cdot | X^q, Z)$
 841 **Ensure:** Final answer Y^a (and optional reasoning chain Y^r)

842
 843
 844
 845
 846
 847
 848
 849
 850
 851
 852
 853
 854
 855
 856
 857
 858
 859
 860
 861
 862
 863