
ImgCoT: Compressing Long Chain of Thought into Compact Visual Tokens for Efficient Reasoning of Large Language Model

Anonymous Authors¹

Abstract

Compressing long chains of thought (CoT) into compact latent tokens is crucial for efficient reasoning with large language models (LLMs). Recent studies employ autoencoders to achieve this by reconstructing textual CoT from latent tokens, thus encoding CoT semantics. However, treating textual CoT as the reconstruction target forces latent tokens to preserve surface-level linguistic features (e.g., word choice and syntax), introducing a strong linguistic inductive bias that prioritizes linguistic form over reasoning structure and limits logical abstraction. Thus, we propose ImgCoT that replaces the reconstruction target from textual CoT to the visual CoT obtained by rendering CoT into images. This substitutes linguistic bias with spatial inductive bias, i.e., a tendency to model spatial layouts of the reasoning steps in visual CoT, enabling latent tokens to better capture global reasoning structure. Moreover, although visual latent tokens encode abstract reasoning structure, they may blur reasoning details. We thus propose a loose ImgCoT, a hybrid reasoning that augments visual latent tokens with a few key textual reasoning steps, selected based on low token log-likelihood. This design allows LLMs to retain both global reasoning structure and fine-grained reasoning details with fewer tokens than the complete CoT. Extensive experiments across multiple datasets and LLMs demonstrate the effectiveness of the two versions of ImgCoT.

1. Introduction

Chain-of-Thought (CoT) (Wei et al., 2022; Chen et al., 2025b) reasoning enhances the reasoning ability of large language models (LLMs) but incurs substantial reasoning over-

¹Anonymous Institution, Anonymous City, Anonymous Region, Anonymous Country. Correspondence to: Anonymous Author <anon.email@domain.com>.

Preliminary work. Under review by the International Conference on Machine Learning (ICML). Do not distribute.

head due to lengthy intermediate traces (Sui et al., 2025).

To address this issue, recent studies propose latent CoT reasoning (Deng et al., 2024; Hao et al., 2024; Shen et al., 2025; Tan et al., 2025; Su et al., 2025), which aims to compress long textual CoT into a small set of latent tokens. Among them, autoencoder-based compression (Su et al., 2025) has shown strong reasoning performance. Typically, as illustrated in the left part of Figure 1, it applies an autoencoder (Hinton & Salakhutdinov, 2006; Vincent et al., 2008; Van Den Oord et al., 2017) framework, where an encoder maps textual CoT into compact latent embeddings, and a decoder is trained to reconstruct the original CoT from these latent tokens. Through this reconstruction objective, latent tokens are expected to encode the semantic information of the CoT, enabling LLMs to reason over compressed representations instead of explicit text.

Despite their success, existing approaches suffer from a fundamental limitation: they treat textual CoT as the reconstruction target. As a result, latent tokens are forced to preserve surface-level linguistic patterns, such as word choice, syntactic structures, and stylistic variations. This introduces a strong linguistic inductive bias, causing latent representations to prioritize language form rather than reasoning structure. The reconstructed outputs clearly reveal this core issue. As illustrated in the left part of Figure 1, although the trained autoencoder can reconstruct a linguistically fluent CoT from textual latent tokens, the recovered CoT often contains structural errors compared with the original CoT. For example, four rules in the original CoT are decoded into only three, with rule 4 omitted. Moreover, the reasoning dependency in rule 2 and rule 3 is incorrectly rewritten, altering the original logical flow. In general, latent tokens with linguistic inductive bias tend to memorize how reasoning is expressed instead of capturing how reasoning is organized, which hinders logical abstraction.

We argue that an effective latent reasoning representation should focus on reasoning structure rather than linguistic realization. To this end, we propose ImgCoT, which replaces textual CoT reconstruction with visual CoT reconstruction. Specifically, as shown in the middle part of Figure 1, we render textual CoT into images that explicitly encode the layout of reasoning steps (e.g., step segmentation, hierarchi-

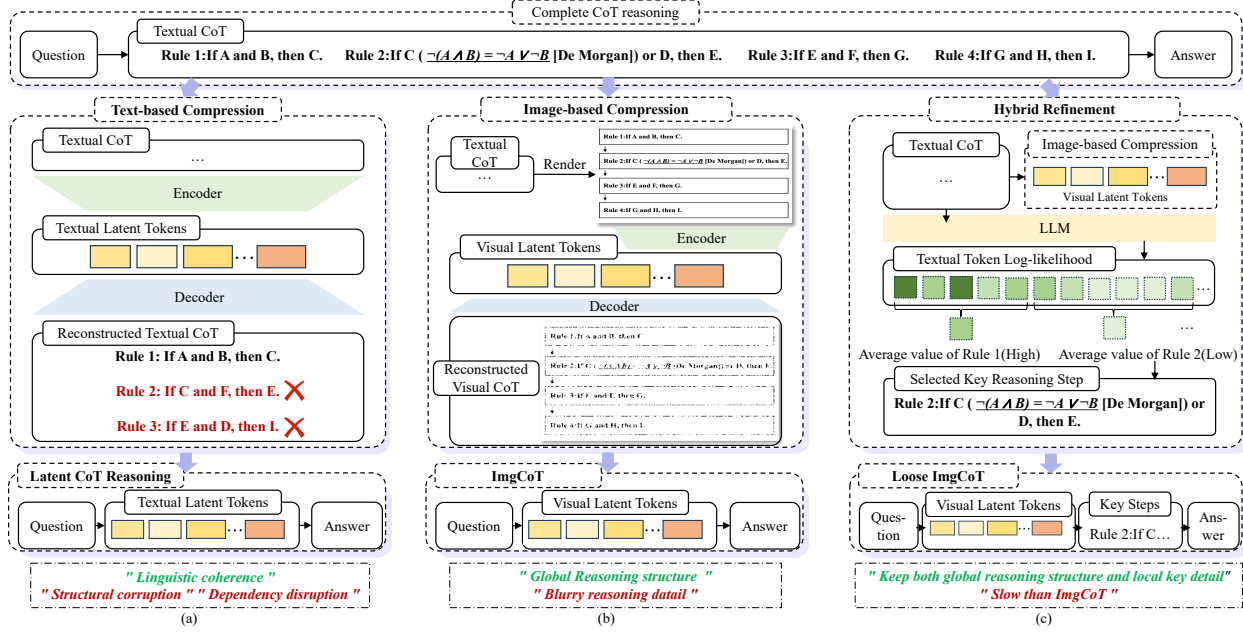


Figure 1. Comparison between three different latent reasoning methods. (a) The reconstructed textual CoT is linguistically fluent but contains structural errors compared with the original CoT, revealing that latent tokens in text-based compression tend to memorize expression patterns rather than reasoning organization. (b) The well-structured CoT in reconstructed visual CoT reflects that latent tokens in image-based compression preserve the global reasoning structure but blur the local reasoning details. (c) Loose ImgCoT further allocates a small number of additional tokens to allow LLMs to explicitly generate critical reasoning steps under uncertainty, thereby combining global reasoning structure with local reasoning precision.

cal structure, and logical flow). Instead of reconstructing text, the autoencoder reconstructs these visual CoT. This design fundamentally changes the inductive bias of the compression process: text reconstruction enforces a linguistic inductive bias, while image reconstruction introduces a spatial inductive bias, encouraging the model to capture global reasoning layouts rather than word-level details. Thus, ImgCoT enables latent tokens to better abstract the high-level reasoning flow, such as step dependencies and logical progression, while discarding unnecessary linguistic variations.

However, from the reconstructed visual CoT in the image-based compression shown in Figure 1, we observe that while visual latent tokens effectively encode global reasoning structures, they may obscure fine-grained reasoning details. For general-purpose knowledge, this blurring has a limited impact, as LLMs already possess strong prior mastery of such skills. However, it can substantially affect domain-specific reasoning skills, where precise operational details are crucial. For example, De Morgan’s law in Rule 2 of Figure 1, a fundamental logical rule for simplifying expressions, is difficult to completely encode into visual latent tokens. This limitation may lead to degraded reasoning performance in logic-related domains.

To mitigate this issue, we further propose Loose ImgCoT (L-ImgCoT), a hybrid reasoning paradigm that augments visual latent tokens with a small number of textual reasoning steps.

These steps are automatically selected based on low token log-likelihood, which identifies uncertain or difficult reasoning steps that the LLMs struggle to infer implicitly. This hybrid design strikes a balance between global abstraction from visual latent tokens and local precision from selected textual steps, allowing LLMs to retain essential reasoning details while still benefiting from substantial token reduction compared to full CoT generation.

Our contributions can be summarized as follows:

- We reveal that visual inductive bias is more suitable than linguistic inductive bias for compressing CoT, as it better preserves global reasoning structure.
- We propose ImgCoT, the first visual CoT compression framework, and L-ImgCoT, an extension that retains fine-grained reasoning details.
- Extensive experiments across 4 datasets and 3 LLM backbones demonstrate the effectiveness of ImgCoT and L-ImgCoT.

2. Related Work

2.1. Latent CoT Reasoning

Compared with other CoT compression methods (Chen et al., 2025a; Kang et al., 2025; Li et al., 2025; Ma et al.,

2025), latent-space CoT reasoning, which avoids generating long textual reasoning traces during inference, has recently attracted significant attention. ICoT (Deng et al., 2024), Coconut (Hao et al., 2024), and CCoT (Cheng & Van Durme, 2024) pioneered the concept of latent reasoning. Specifically, ICoT leverages curriculum learning to gradually internalize explicit textual CoTs into the forward propagation of LLMs, while Coconut and CCoT compress long textual CoTs into continuous latent vectors via curriculum learning. Subsequent works, such as CODI (Shen et al., 2025) and SynAdapt (Wang et al., 2025), further distill the reasoning capability of complete CoT-supervised models into latent-reasoning LLMs, improving latent reasoning performance. Beyond curriculum learning, CoLaR (Tan et al., 2025) compresses CoTs by pooling token representations across those of textual CoT, while LightThinker (Zhang et al., 2025) employs attention mechanisms for CoT compression.

However, although these methods accelerate CoT reasoning, their performance still falls short of full textual CoT reasoning. Recently, Su et al. (Su et al., 2025) broke this limitation by achieving latent reasoning without sacrificing performance. Specifically, they adopt Vector-Quantized Variational Autoencoders (VQ-VAEs) (Van Den Oord et al., 2017; Razavi et al., 2019) to compress the prefix of full textual CoTs, and train LLMs to reason jointly over compressed latent tokens and remaining textual steps. Nevertheless, as we analyze before, linguistic inductive bias limits the preservation of abstract reasoning structures in latent tokens, leading to suboptimal performance. Moreover, since most textual CoT steps still participate in inference, the acceleration gain remains limited.

2.2. Image Tokenization

Learning compact representations for images (i.e., image tokenization) has long been a central topic in generative modeling. Early studies attempt the autoencoder (AE) (Hinton & Salakhutdinov, 2006; Vincent et al., 2008) paradigm, which learns to map high-dimensional images into compact latent representations and reconstruct them back. Building upon this foundation, Variational Autoencoders (VAEs) (Kingma & Welling, 2013) introduce probabilistic latent variables, enabling continuous distribution modeling and facilitating downstream generative frameworks. To further discretize the latent space and ease sequence modeling, VQ-VAEs (Van Den Oord et al., 2017; Razavi et al., 2019) replace continuous latents with codebook-based discrete tokens, establishing a cornerstone for modern image tokenization.

Subsequent works extend these foundations along multiple dimensions, including adversarial (Esser et al., 2021; Chang et al., 2022) and perceptual (Zhang et al., 2018) training objectives for improving visual fidelity, transformer-based tokenizers (Yu et al., 2021; Cao et al., 2023) for enhanced mod-

eling capacity, and advanced quantization strategies such as multi-stage (Lee et al., 2022; Zheng et al., 2022) or lookup-free (Yu et al., 2023b; Mentzer et al., 2023) discretization. Despite these improvements, most existing methods inherit a common design choice: encoding images into a 2D grid of local tokens, which biases each token toward local spatial content.

Recently, 1D tokenization (Yu et al., 2024; Beyer et al., 2025) has been proposed to learn a latent sequence in which each token captures global spatial information, enabling significantly higher compression ratios by exploiting long-range redundancy. In our work, to preserve global reasoning structure within latent tokens, we adopt a 1D tokenization model, TiTok (Yu et al., 2024), as the backbone for compressing visual CoTs.

3. Method

ImgCoT consists of three stages: (1) Visual text tokenization, (2) LLM training with visual latent tokens, and (3) LLMs inference in latent space. Furthermore, we introduce Loose ImgCoT, a hybrid reasoning strategy that augments visual latent tokens with selectively preserved textual steps to retain fine-grained reasoning details.

3.1. ImgCoT

3.1.1. VISUAL TEXT TOKENIZATION

Given a text sequence, we first render it into a visual text image $I \in \mathbb{R}^{H \times W \times 3}$. Specifically, as shown in Figure 2, we segment the input text based on predefined delimiters (e.g., “\n”), wrap each segment with a bounding box, and annotate inter-segment dependencies using directed arrows. This rendering converts linguistic input into a spatial representation, enabling subsequent modeling to exploit structural layouts rather than surface linguistic forms.

We adopt TiTok (Yu et al., 2024) as the autoencoder to compress I into discrete latent tokens. First, the visual encoder \mathcal{E} maps the I into continuous feature embeddings:

$$h = \{h_i\}_{i=0}^n = \mathcal{E}(I), h_i \in \mathbb{R}^d. \quad (1)$$

To obtain discrete representations, we maintain a learnable codebook:

$$\mathcal{C} = \{e_j\}_{j=0}^k, e_j \in \mathbb{R}^d. \quad (2)$$

Then, each latent embedding h_i is quantized by nearest-neighbor lookup in the \mathcal{C} , i.e.,

$$z_i = \operatorname{argmin}_{e_j \in \mathcal{C}} \|h_i - e_j\|_2. \quad (3)$$

This process converts continuous embeddings into discrete visual tokens, forcing the model to express I using a finite set of reusable visual prototypes. Since the quantization operation is non-differentiable, we adopt the straight-through

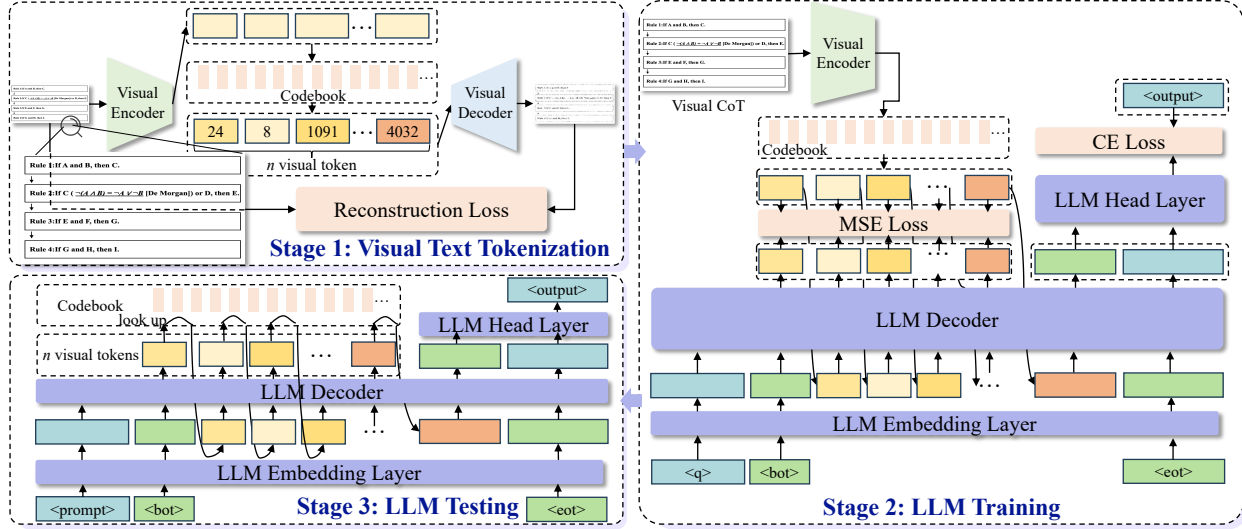


Figure 2. The overall training pipeline of ImgCoT. First, a visual encoder is trained by reconstructing visualized text, thereby compressing textual inputs into latent tokens. Then, the LLM is trained in an autoregressive manner to perform reasoning over the visual latent tokens. $\langle \cdot \rangle$ denotes the output sequence processed by the LLM tokenizer.

estimator to enable end-to-end training:

$$\hat{z}_i = h_i + sg(z_i - h_i), \quad (4)$$

where sg denotes stop-gradient. The \hat{z}_i are then fed into the visual decoder \mathcal{D} to obtain reconstruct result $\hat{I} = \mathcal{D}(\hat{z})$. Finally, the reconstruction loss

$$\mathcal{L}_{rec} = \|\hat{I} - I\|_2^2, \quad (5)$$

is applied to encourage the latent tokens \hat{z}_i to capture the key information, such as the structural layout of the input text. Notably, this subsection only outlines the high-level framework, and further implementation details can be found in the TiTok (Yu et al., 2024).

3.1.2. LLM TRAINING WITH VISUAL LATENT TOKENS

The LLMs training process is shown in stage 2 of Figure 2. Specifically, let q denote the question, c the corresponding textual CoT, and a the final answer. We first apply the visual encoder \mathcal{E} and codebook \mathcal{C} learned in the first stage to compress c into z , which consists of n latent tokens. Meanwhile, we tokenize q and a using the LLM tokenizer, obtaining the token ids $\langle q \rangle = \{q_i\}_{i=0}^{|\langle q \rangle|}$ and $\langle a \rangle = \{a_i\}_{i=0}^{|\langle a \rangle|}$, respectively, where $\langle \cdot \rangle$ is the output sequence processed by the LLM tokenizer and $|\cdot|$ denotes the length of the input. Then, a training sample is constructed as

$$x = \{\langle q \rangle, \langle bot \rangle, z, \langle eot \rangle, \langle output \rangle\} = \{x_i\}_{i=0}^{|x|}, \quad (6)$$

where $\langle output \rangle = \langle a \rangle \oplus \langle eot \rangle$, $\langle bot \rangle \langle eot \rangle$ are the special token ids and \oplus is the concatenation operation. Subsequently, we train the LLM in a standard autoregressive manner to

perform reasoning over latent tokens and generate the final answer, formulated as follows:

$$\mathcal{L}_{ar} = \frac{\sum_{i=s_z}^{e_z} \mathcal{L}_{mse}(f_\theta(x_{:i}), x_i) + \sum_{i=s_o}^{e_o} \mathcal{L}_{ce}(f_\theta(x_{:i}), x_i)}{n + |\langle output \rangle|}, \quad (7)$$

where s_z and e_z denote the start and end indices of z in the x , and s_o and e_o denote those of $\langle output \rangle$ in x . $f_\theta(x_{:i})$ represents the LLM output at the position $i - 1$ given the input sequence $x_{:i}$, parameterized by θ . Here, $x_{:i}$ denotes the subsequence of x spanning indices from 0 to $i - 1$. \mathcal{L}_{mse} and \mathcal{L}_{ce} denote the mean square error and cross-entropy loss, respectively.

Through this training procedure, the LLM learns to perform reasoning directly in the visual latent space. Since n is significantly smaller than $|\langle c \rangle|$, the LLM can obtain the answer more efficiently than standard CoT reasoning.

3.1.3. TESTING PROCEDURE

As shown in stage 3 of Figure 2, inference is performed in an autoregressive manner. Importantly, the visual encoder is discarded at test time, ensuring that our method incurs no extra computational cost in real-world applications.

3.2. Loose ImgCoT

While visual latent tokens effectively capture global reasoning structure, they may obscure fine-grained domain-specific reasoning skills, thereby hindering LLMs from acquiring these skills during fine-tuning. To address this limitation, we propose L-ImgCoT, which augments visual latent tokens with a small set of critical textual reasoning steps.

L-ImgCoT follows the same training and inference procedures as ImgCoT. The only difference lies in the construction of training data. Following prior studies (Wan et al., 2024; Anonymous, 2023), the average log-likelihood of tokens in an input sequence reflects the LLM’s confidence in understanding this input. Therefore, before fine-tuning, we compute the average token-level log-likelihood of the general pretraining corpus, denoted as γ , which represents the understanding level of LLMs to the general knowledge. Formally, the token-level log-likelihood at the position i of an input text sequence t can be evaluated as:

$$\text{conf}(t, i) = \log(\text{softmax}(f_{\theta}(t_{:i}))^{(t_i)}), \quad (8)$$

where $\text{softmax}(f_{\theta}(t_{:i}))^{(t_i)}$ denotes the probability predicted by LLMs that the token at position i is $\langle t_i \rangle$. Based on it, we can get γ by

$$\gamma = \sum_{t^j \in \mathcal{T}} \sum_{i=0}^{|t^j|} \text{conf}(t^j, i), \quad (9)$$

where $\mathcal{T} = \{t^j\}_{j=0}^{|\mathcal{T}|}$ is the general pretraining corpus.

Then, we use γ as a threshold to filter each reasoning step c^i in $c = \{c^i\}_{i=0}^{|c|}$. Specifically, a reasoning step c^j is filtered out and be replaced with “...” if its average log-likelihood is greater than γ . Note that multiple consecutive filtered steps are replaced by a single ellipsis. After filtering, we obtain a refined reasoning trace $\hat{c} = \{\hat{c}^i\}_{i=0}^{|\hat{c}|}$. Here, $|\langle \hat{c} \rangle| < |\langle c \rangle|$.

As a result, we can replace the $\langle \text{output} \rangle$ in Equation (6) with $\langle \text{output} \rangle = \langle \hat{c} \rangle \oplus \langle a \rangle \oplus \langle \text{eot} \rangle$ to get the training data for L-ImgCoT. With this training data construction strategy, L-ImgCoT enables LLMs to reason primarily in the latent space for global structural abstraction, while selectively incorporating a small number of textual tokens to preserve domain-specific reasoning skills. This hybrid design achieves a favorable trade-off between reasoning efficiency and fine-grained logical precision.

4. Experiments

In this section, we first present the experimental setup in detail. We then systematically address the following key questions through the carefully designed experimental tasks:

- **RQ1:** How do our methods perform across different datasets and LLMs?
- **RQ2:** Why does image-based compression work better than text-based compression?
- **RQ3:** What other advantages does spatial inductive bias offer over linguistic inductive bias?
- **RQ4:** Does L-ImgCoT effectively preserve critical fine-grained reasoning skills during inference?

4.1. Experimental Setting

4.1.1. BENCHMARKS

We evaluate the proposed methods across three representative reasoning scenarios: mathematical reasoning, commonsense reasoning, and logical reasoning. For mathematical reasoning, we adopt GSM8K (Cobbe et al., 2021) and MATH (Hendrycks et al., 2021) as benchmarks. For commonsense reasoning, we evaluate our methods on GPQA-extended (Rein et al., 2024). For Logical reasoning, we use ProsQA (Hao et al., 2024) to assess logical reasoning performance. More details of these dataset can be found in the Section A.1.

4.1.2. IMPLEMENTATION

Several key implementation details are below. More comprehensive descriptions are provided in the Section A.2.

1) Visual Text Tokenization. We train the autoencoder using MathPile (Wang et al., 2024) as the training corpus. The procedure for rendering text into images follows the implementation of Glyph (Cheng et al., 2025), as it provides the most suitable hyperparameters for rendering. The image resolution $H \times W$ is fixed to 512×512 . Unless otherwise specified, we set the number of latent tokens to $n = 8$, the codebook size to $k = 4096$, and the latent dimension d to be equal to the hidden size of the downstream LLM.

2) LLM Fine-tuning. Due to computational constraints, we adopt different fine-tuning strategies based on model size. For models with fewer than 1B parameters, we perform full-parameter fine-tuning. Otherwise, we apply LoRA (Hu et al., 2022) ($r = 12, \alpha = 32$) to conduct efficient fine-tuning.

4.2. RQ1 (Main Results)

4.2.1. COMPARISON WITH OTHER BASELINE METHODS

We compare ImgCoT and L-ImgCoT with the baseline Full-CoT and several recent, reproducible state-of-the-art (SOTA) latent reasoning methods, including Coconut (Hao et al., 2024), ICoT (Deng et al., 2024), CODI (Shen et al., 2025), and CoLaR (Tan et al., 2025), where Full-CoT supervises LLMs using complete textual CoT. The comparison results are summarized in Table 1. Note that, we implement Full-CoT ourselves, while all other SOTA methods are reproduced using the official implementation of CoLaR.

We observe that: (1) compared to other SOTA latent reasoning methods, ImgCoT achieves better and more stable performance under comparable reasoning costs. (2) ImgCoT matches or even outperforms Full-CoT in many cases, highlighting the importance of abstracting correct reasoning structures during reasoning. In contrast, Full-CoT’s token-by-token generation of complete reasoning traces is often

Table 1. Comparison of accuracy (*Acc*) and token consumption of reasoning (*# Tokens*) across different methods. *w/* textual tokens denotes textual latent reasoning based on text compression, while *w/o* layout indicates directly rendering text into images without explicit logical dependency symbols (e.g., arrows). "-" represents that the result is missing due to the limitation of computational resources. Bold and underlined numbers indicate the best and second-best performance, respectively.

Model	MATH		GSM		GPQA		ProsQA		
	<i>Acc</i> ↑	<i># Tokens</i> ↓	<i>Acc</i> ↑	<i># Tokens</i> ↓	<i>Acc</i> ↑	<i># Tokens</i> ↓	<i>Acc</i> ↑	<i># Tokens</i> ↓	
Qwen2.5-0.5B -Instruction	Full-CoT	9.2	149.4	<u>16.9</u>	116.3	34.5	150.2	94.6	71.3
	Coconut (COLM'25)	3.5	6.0	3.7	6.0	26.0	6.0	73.8	6.0
	ICoT	3.3	0.0	3.9	0.0	39.6	0.0	54.9	0.0
	CODI (EMNLP'25)	0.4	6.0	0.5	6.0	27.1	6.0	0.0	6.0
	CoLaR (NeurIPS'25)	3.1	39.4	5.8	19.1	27.7	28.7	68.9	18.9
	ImgCoT(ours)	<u>9.8</u>	8.0	9.2	8.0	34.5	8.0	<u>97.4</u>	8.0
	<i>w/o</i> layout	9.1	8.0	9.2	8.0	32.7	8.0	96.8	8.0
	<i>w/</i> textual tokens	9.0	8.0	8.3	8.0	27.3	8.0	96.2	8.0
	L-ImgCoT(ours)	10.1	102.8	17.5	64.7	<u>38.1</u>	89.3	98.6	40.9
Qwen2.5-1.5B -Instruction	Full-CoT	19.4	238.7	<u>44.1</u>	126.6	40.0	229.5	99.6	46.2
	Coconut (COLM'25)	8.8	6.0	5.4	6.0	31.4	6.0	<u>99.6</u>	6.0
	ICoT	12.5	0.0	22.3	0.0	30.7	0.0	98.6	0.0
	CODI (EMNLP'25)	4.3	6.0	3.6	6.0	<u>42.8</u>	6.0	<u>99.6</u>	6.0
	CoLaR (NeurIPS'25)	4.9	60.5	6.9	22.9	37.1	30.8	99.4	7.1
	ImgCoT(ours)	<u>19.5</u>	8.0	38.7	8.0	41.8	8.0	100.0	8.0
	<i>w/o</i> layout	19.1	8.0	38.4	8.0	36.4	8.0	100.0	8.0
	<i>w/</i> textual tokens	17.4	8.0	37.5	8.0	36.4	8.0	99.2	8.0
	L-ImgCoT(ours)	19.9	127.4	45.2	71.3	43.6	130.8	100.0	21.7
LLama3.2-3B -Instruction	Full-CoT	23.8	205.9	<u>60.5</u>	105.6	45.5	179.3	100.0	41.3
	Coconut (COLM'25)	13.8	6.0	19.5	6.0	28.6	6.0	100.0	6.0
	ICoT	15.0	0.0	46.7	0.0	28.9	0.0	<u>99.2</u>	0.0
	CODI (EMNLP'25)	-	-	23.3	6.0	-	-	-	-
	CoLaR (NeurIPS'25)	19.1	41.2	27.4	20.1	20.4	33.9	100.0	6.8
	ImgCoT(ours)	<u>24.1</u>	8.0	56.8	8.0	<u>43.6</u>	8.0	100.0	8.0
	<i>w/o</i> layout	23.6	8.0	52.7	8.0	43.6	8.0	100.0	8.0
	<i>w/</i> textual tokens	22.6	8.0	49.3	8.0	41.8	8.0	100.0	8.0
	L-ImgCoT(ours)	24.3	143.8	61.0	63.2	45.5	115.7	100.0	27.5

vulnerable to minor local errors that propagate to incorrect final answers. (3) L-ImgCoT consistently outperforms Full-CoT across all settings while reducing inference cost by approximately 30%, mitigating ImgCoT's tendency to blur critical reasoning details and offering a practical alternative for scenarios where inference latency is less constrained.

4.2.2. ABLATION STUDIES

In Table Table 1, we present two key ablation studies. First, we compare visual latent reasoning (ImgCoT) with text-based latent reasoning (*w/* textual tokens) in terms of performance. The text-based latent reasoning model shares the same architecture and training strategy as ImgCoT, differing only in that its reconstruction target is textual CoT, which replaces the reconstruction loss in Equation (5) with the next-token prediction loss. The results show that visual latent reasoning consistently outperforms its textual counterpart, suggesting that spatial inductive bias is more suitable than linguistic inductive bias for CoT compression.

Second, we examine the impact of directly rendering textual reasoning steps into images (*w/o* layout) versus augmenting renderings with explicit symbolic cues (e.g., arrows) that encode reasoning logic (ImgCoT). We find that direct rendering yields slightly inferior reasoning performance compared to ImgCoT. This demonstrates the importance of accurately modeling logical dependencies among reasoning steps in the visual space, and further suggests that visual latent tokens indeed preserve structural dependencies between reasoning steps, as otherwise the quality of spatial dependency modeling for reasoning steps would not affect reasoning performance.

4.3. RQ2 (Image-based Compression vs. Text-based Compression)

4.3.1. QUALITATIVE COMPARISON

We first provide a qualitative comparison. As shown in Figure 3, we visualize reconstructions decoded from latent tokens under text-based and image-based compression. We

Original CoT	Reconstructed Textual CoT	Reconstructed Visual CoT
<p>If $Sb=-1, d=5S$, then $Sac^2+bd=ac^2-2S$, so $Sac^2=6S$. We substitute $Sa=1-cS$ from the first equation to get the quadratic $S^2c^2-c-6=0S$, which has solutions $S^2c=2S$ (so $Sa=3S$) or $S^2c=3S$ (so $Sa=2S$). In either case, we get that $Sa+b+c+d=boxed{5S}$.</p> <p>The remaining equation, $Sad + bc = 17S$, tells us that the coefficients are $Sa = 3, b = -1, c = -2, d = 5S$.</p> <p>Since the numbers less than 15 are small, we can easily apply this process to all the numbers from 1 to 15. Here is a table showing how many factors each number has:</p> <pre> \begin{tabular}{ c c } \hline number & how many factors \\ \hline 1 & 1 \\ 2 & 2 \\ 3 & 2 \\ 4 & 3 \\ 5 & 2 \\ 6 & 4 \\ 7 & 2 \\ 8 & 4 \\ 9 & 3 \\ 10 & 4 \\ 11 & 2 \\ 12 & 6 \\ 13 & 2 \\ 14 & 4 \\ \hline \end{tabular} </pre>	<p>Let SAS be a positive integer and SBS a positive integer. Let $SA^2=!(1,2,3,4,5)S$ and $SB^2=!(1,2,3,4,5)S$.</p> <p>Then $SA^2=!(1,2,3,4,5)S$ and $SB^2=!(1,2,3,4,5)S$.</p> <p>I know that $SA^2=!(1,2,3,4,5)S$ and $SB^2=!(1,2,3,4,5)S$.</p>	<p>If $d=-1, b=156$, that does not work. If $d=25$, so $Sa=3S$, we have $Sa^2=9S^2, b^2=225S^2$, and $Sa^2+b^2=234S^2$. We want $Sa^2+b^2=234S^2$, so $Sa=3S, b=156$. In either case, we see that $Sa+b+c+d=boxed{5S}$.</p> <p>The same rare case, $ad = bc = 17S$, tells us that the coefficients are $Sa = 3, b = -1, c = -2, d = 5S$.</p> <pre> \begin{tabular}{ c c } \hline length (number of bits) & position in position, word size, trace \\ \hline 1 & 1 \\ 2 & 2 \\ 3 & 3 \\ 4 & 4 \\ 5 & 5 \\ 6 & 6 \\ 7 & 7 \\ 8 & 8 \\ 9 & 9 \\ 10 & 10 \\ 11 & 11 \\ 12 & 12 \\ 13 & 13 \\ 14 & 14 \\ 15 & 15 \\ \hline \end{tabular} </pre>

Figure 3. Qualitative comparison between the original CoT and its reconstructions from different latent representations. Text-based compression preserves surface-level linguistic forms, whereas image-based compression captures the global reasoning skeleton and stepwise layout of CoT. These examples illustrate the intrinsic inductive biases of different latent representations.

observe that (1) reconstructed textual CoTs remain linguistically fluent but exhibit weak semantic alignment with the original CoTs, indicating that text-based compression fails to preserve the underlying semantics of the original CoT. In contrast, reconstructed visual CoTs preserve both the logical dependencies across reasoning steps and the internal logic within individual steps, for example, faithfully recovering the direction of arrows in both cases and the exact table row count in the second reasoning step of the case shown in the second row. These reconstruction differences directly reflect the distinct types of information retained in latent tokens by text-based versus image-based compression: the former tends to encode surface-level linguistic form, whereas the latter prioritizes abstractions of reasoning structure. Since accurate modeling of reasoning structure is more critical for reasoning performance, image-based compression is therefore more effective for CoT compression.

4.3.2. QUANTITATIVE COMPARISON

We next present a quantitative comparison. As shown in Figure 4, we evaluate LLM reasoning performance as the number of latent tokens n is gradually reduced under the two compression strategies. As n decreases, latent tokens retain progressively less information, requiring each compression method to prioritize which information is most critical to retain for effective reconstruction. We observe that as n decreases, the performance gap between text-based and image-based compression steadily widens. This indicates that the information prioritized by image-based compression in latent tokens is more beneficial for LLM reasoning than that preserved by text-based compression, i.e., spatial inductive bias is better than linguistic inductive bias. We observe that model performance degrades when n becomes very large. Due to space limitations, we defer a detailed




Figure 4. Performance comparison under varying numbers of visual versus textual latent tokens. The LLM is Qwen2.5-0.5B-Instruction.

analysis of this phenomenon to Section C.

4.4. RQ3 (Generalization of Spatial Inductive Bias)

In this subsection, we further demonstrate a key advantage of prioritizing structural dependencies in latent tokens (spatial inductive bias) over preserving surface-level lexical compositions (linguistic inductive bias): generalization. We argue that reasoning structures capture abstract logical patterns, which are more broadly transferable across reasoning tasks than task-specific lexical realizations. Accordingly, we compare the generalization ability of spatial and linguistic inductive biases, as shown in Table 2.

Specifically, we fine-tune LLMs on MetaMathQA (Yu et al., 2023a), while compressing its CoT using text-based or image-based compression, respectively. MetaMathQA is an augmented combination of the MATH and GSM datasets, and we therefore evaluate the resulting models on out-of-domain benchmarks, including Gaokao-Math-2023 (Liao et al., 2024), Svamp (Patel et al., 2021), MultiArith (Roy & Roth, 2015) and SingleEq (Koncel-Kedziorski et al., 2015). As shown in the table, visual latent reasoning with spatial inductive bias substantially outperforms text-based latent reasoning with linguistic inductive bias on all out-of-domain benchmarks, demonstrating the strong generalization capability of spatial inductive bias in reasoning tasks.

4.5. RQ4 (Fine-grained Reasoning with L-ImgCoT)

In this section, we first demonstrate the effectiveness of the proposed strategy for retaining key reasoning steps. As illustrated in Figure 5 (additional examples are provided in the Section B), we observe that the retained reasoning steps correspond to essential domain-specific reasoning skills, such as writing code to generate plots (blue example) and applying mathematical theories (yellow example). These skills

Table 2. Out-of-domain generalization of visual and textual latent tokens at inference. The LLM is Qwen2.5-0.5B-Instruction.

Model	In-Domain		Out-of-Domain		
	GSM	Gaokao	SVAMP	SingleEq	MultiArith
ImgCoT	11.4	3.1	9.6	12.0	5.8
w/o visual tokens	10.0	1.8	5.6	8.4	4.2
Difference	↓1.4	↓1.3	↓3.0	↓3.6	↓1.6

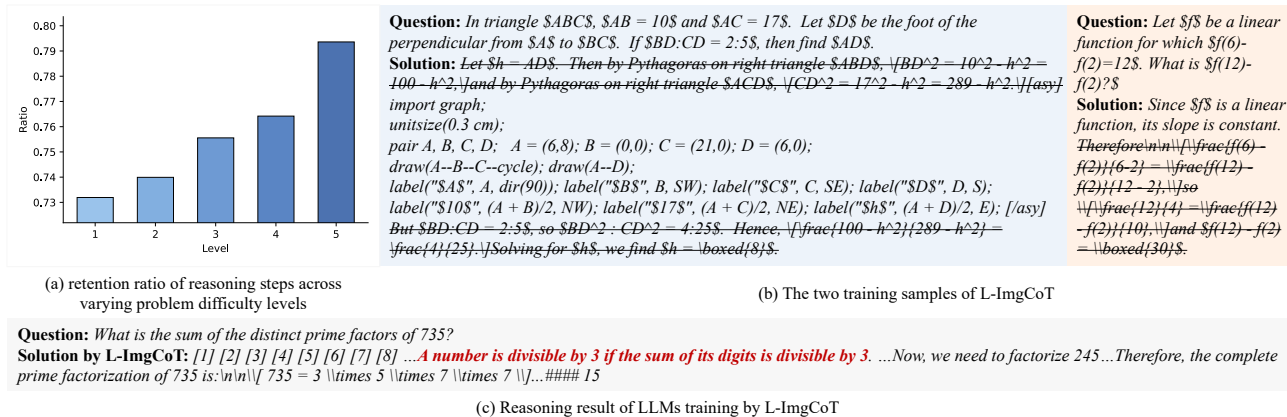


Figure 5. Effectiveness of Retaining Critical Reasoning Steps. The strikethrough text in (b) corresponds to reasoning deemed unimportant and filtered out by our strategy. [·] represents the latent tokens. The LLM adopted here is Qwen2.5-0.5B-Instruction.

Table 3. Performance Comparison of L-ImgCoT With and Without Visual Latent Tokens. The LLM is Qwen2.5-0.5B-Instruction.

Method	MATH	GSM	GPQA	ProsQA
L-ImgCoT	9.1	15.8	29.1	94.2
w/o visual tokens	10.1	17.5	38.1	98.6

are critical for general-purpose LLMs to acquire, thereby validating the effectiveness of the proposed key-step retention strategy. In addition, we analyze the proportion of retained reasoning steps across different difficulty levels on the MATH dataset. More difficult problems typically involve a larger number of critical reasoning steps, and the retention ratios shown in the figure reflect this trend, supporting the rationality of the proposed retention strategy. Beyond training examples, the figure also presents the actual reasoning behavior of LLMs after L-ImgCoT training. We observe that the fine-tuned LLMs learn to retain only key reasoning steps, for example, outputting only empirical laws (red text) in mathematics.

Second, we demonstrate the importance of inserting visual latent tokens before key reasoning steps. As shown in the table, we compare LLM performance with and without visual latent tokens. We observe that LLMs perform substantially worse without visual latent tokens than when they are included. We attribute this difference to the fact that visual

latent tokens encode global reasoning structure, which plays a critical role in guiding subsequent reasoning. Analogously, when humans form a global reasoning plan in advance, many low-level details can be omitted, and only the outcomes of key reasoning steps need to be made explicit. In contrast, without such global abstraction, each individual reasoning step becomes necessary and cannot be safely omitted.

5. Conclusion

In this work, we revisit CoT compression through the lens of inductive bias and show that what is preserved in latent representations matters more than how much information is retained. Our key finding is that visual inductive bias, which prioritizes global reasoning structure over surface-level linguistic form, is fundamentally more suitable than linguistic bias for compressing CoT. Based on it, we propose ImgCoT, the framework that compresses CoT into visual latent tokens, and L-ImgCoT, an extension that further retains fine-grained reasoning details under lower inference budgets. Extensive experiments across multiple datasets and LLM backbones demonstrate the effectiveness of our proposed methods.

Looking forward, our work opens several directions, including more expressive structure-aware rendering schemes, richer visual elements beyond textual CoT, and more stable training paradigms. We hope this work encourages future research to explore visual abstraction as a principled path toward efficient and generalizable reasoning.

Impact Statement

This paper presents work whose goal is to advance the field of Machine Learning by studying more efficient and structured representations for CoT reasoning in LLMs. By examining the role of inductive bias in reasoning compression, our work contributes to improved efficiency and generalization of reasoning systems, which may help reduce reasoning cost and resource consumption in practical deployments.

There are many potential societal consequences of advances in efficient reasoning models, both positive and negative, which are well established in the literature on LLMs. We do not believe that this work introduces new ethical considerations beyond those commonly associated with improving model efficiency and reasoning capability, and therefore do not highlight specific societal impacts here.

References

Anonymous. When less is more: Investigating data pruning for pretraining LLMs at scale. In *NeurIPS Workshop on Attributing Model Behavior at Scale*, 2023. URL <https://openreview.net/forum?id=XUIYn3jo5T>.

Beyer, L. L., Li, T., Chen, X., Karaman, S., and He, K. Highly compressed tokenizer can generate without training. *arXiv preprint arXiv:2506.08257*, 2025.

Cao, S., Yin, Y., Huang, L., Liu, Y., Zhao, X., Zhao, D., and Huang, K. Efficient-vqgan: Towards high-resolution image generation with efficient vision transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 7368–7377, 2023.

Chang, H., Zhang, H., Jiang, L., Liu, C., and Freeman, W. T. Maskgit: Masked generative image transformer. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 11315–11325, 2022.

Chen, X., Zhou, S., Liang, K., Sun, X., and Liu, X. Skip-thinking: Chunk-wise chain-of-thought distillation enable smaller language models to reason better and faster. In Christodoulopoulos, C., Chakraborty, T., Rose, C., and Peng, V. (eds.), *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pp. 12142–12157, Suzhou, China, November 2025a. Association for Computational Linguistics. ISBN 979-8-89176-332-6. doi: 10.18653/v1/2025.emnlp-main.610. URL <https://aclanthology.org/2025.emnlp-main.610/>.

Chen, X., Zhou, S., Liang, K., Yuan, D., Chen, H., Sun, X., Meng, L., and Liu, X. Putting on the thinking hats: A survey on chain of thought fine-tuning from the perspective of human reasoning mechanism, 2025b. URL <https://arxiv.org/abs/2510.13170>.

Cheng, J. and Van Durme, B. Compressed chain of thought: Efficient reasoning through dense representations. *arXiv preprint arXiv:2412.13171*, 2024.

Cheng, J., Liu, Y., Zhang, X., Fei, Y., Hong, W., Lyu, R., Wang, W., Su, Z., Gu, X., Liu, X., Bai, Y., Tang, J., Wang, H., and Huang, M. Glyph: Scaling context windows via visual-text compression. *arXiv preprint arXiv:2510.17800*, 2025.

Cobbe, K., Kosaraju, V., Bavarian, M., Chen, M., Jun, H., Kaiser, L., Plappert, M., Tworek, J., Hilton, J., Nakano, R., et al. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*, 2021.

Deng, Y., Choi, Y., and Shieber, S. From explicit cot to implicit cot: Learning to internalize cot step by step. *arXiv preprint arXiv:2405.14838*, 2024.

Esser, P., Rombach, R., and Ommer, B. Taming transformers for high-resolution image synthesis. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 12873–12883, 2021.

Hao, S., Sukhbaatar, S., Su, D., Li, X., Hu, Z., Weston, J., and Tian, Y. Training large language models to reason in a continuous latent space. *arXiv preprint arXiv:2412.06769*, 2024.

Hendrycks, D., Burns, C., Kadavath, S., Arora, A., Basart, S., Tang, E., Song, D., and Steinhardt, J. Measuring mathematical problem solving with the math dataset. *NeurIPS*, 2021.

Hinton, G. E. and Salakhutdinov, R. R. Reducing the dimensionality of data with neural networks. *science*, 313 (5786):504–507, 2006.

Hu, E. J., yelong shen, Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., and Chen, W. LoRA: Low-rank adaptation of large language models. In *International Conference on Learning Representations*, 2022. URL <https://openreview.net/forum?id=nZeVKeeFYf9>.

Kang, Y., Sun, X., Chen, L., and Zou, W. C3ot: generating shorter chain-of-thought without compromising effectiveness. In *Proceedings of the Thirty-Ninth AAAI Conference on Artificial Intelligence and Thirty-Seventh Conference on Innovative Applications of Artificial Intelligence and Fifteenth Symposium on Educational Advances in Artificial Intelligence*, AAAI’25/IAAI’25/EAAI’25. AAAI Press, 2025. ISBN 978-1-57735-897-8. doi: 10.1609/aaai.v39i23.34608. URL <https://doi.org/10.1609/aaai.v39i23.34608>.

Kingma, D. P. and Welling, M. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.

- 495 Koncel-Kedziorski, R., Hajishirzi, H., Sabharwal, A., Et-
 496 zioni, O., and Ang, S. D. Parsing algebraic word prob-
 497 lems into equations. *Transactions of the Association for*
 498 *Computational Linguistics*, 3:585–597, 2015. doi: 10.
 499 1162/tacl.a.00160. URL <https://aclanthology.org/Q15-1042/>.
- 500
 501 Lee, D., Kim, C., Kim, S., Cho, M., and Han, W.-S. Autore-
 502 gressive image generation using residual quantization. In
 503 *Proceedings of the IEEE/CVF conference on computer*
 504 *vision and pattern recognition*, pp. 11523–11532, 2022.
- 505
 506 Li, Z., Zhong, J., Zheng, Z., Wen, X., Xu, Z., Cheng, Y.,
 507 Zhang, F., and Xu, Q. Compressing chain-of-thought in
 508 llms via step entropy, 2025. URL <https://arxiv.org/abs/2508.03346>.
- 509
 510 Liao, M., Li, C., Luo, W., Jing, W., and Fan, K. MARIO: MATH reasoning with code interpreter output - a reproducible pipeline. In Ku, L.-W., Martins, A., and Srikumar, V. (eds.), *Findings of the Association for Computational Linguistics: ACL 2024*, pp. 905–924, Bangkok, Thailand, August 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.findings-acl.53. URL <https://aclanthology.org/2024.findings-acl.53/>.
- 511
 512 Ma, X., Wan, G., Yu, R., Fang, G., and Wang, X. CoT-valve:
 513 Length-compressible chain-of-thought tuning. In Che, W.,
 514 Nabende, J., Shutova, E., and Pilehvar, M. T. (eds.), *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 6025–6035, Vienna, Austria, July 2025. Association for Computational Linguistics. ISBN 979-8-89176-251-0. doi: 10.18653/v1/2025.acl-long.300. URL <https://aclanthology.org/2025.acl-long.300/>.
- 515
 516 Mentzer, F., Minnen, D., Agustsson, E., and Tschannen, M. Finite scalar quantization: Vq-vae made simple. *arXiv preprint arXiv:2309.15505*, 2023.
- 517
 518 Patel, A., Bhattamishra, S., and Goyal, N. Are NLP
 519 models really able to solve simple math word prob-
 520 lems? In Toutanova, K., Rumshisky, A., Zettlemoyer,
 521 L., Hakkani-Tur, D., Beltagy, I., Bethard, S., Cot-
 522 terrell, R., Chakraborty, T., and Zhou, Y. (eds.), *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pp. 2080–2094, Online, June 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.naacl-main.168. URL <https://aclanthology.org/2021.naacl-main.168/>.
- 523
 524 Razavi, A., Van den Oord, A., and Vinyals, O. Generating
 525 diverse high-fidelity images with vq-vae-2. *Advances in*
 526 *neural information processing systems*, 32, 2019.
- 527
 528 Rein, D., Hou, B. L., Stickland, A. C., Petty, J., Pang, R. Y.,
 529 Dirani, J., Michael, J., and Bowman, S. R. GPQA: A
 530 graduate-level google-proof q&a benchmark. In *First*
 531 *Conference on Language Modeling*, 2024. URL <https://openreview.net/forum?id=Ti67584b98>.
- 532
 533 Roy, S. and Roth, D. Solving general arithmetic word
 534 problems. In Márquez, L., Callison-Burch, C., and Su,
 535 J. (eds.), *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pp. 1743–1752, Lisbon, Portugal, September 2015. Association for Computational Linguistics. doi: 10.18653/v1/D15-1202. URL <https://aclanthology.org/D15-1202/>.
- 536
 537 Shen, Z., Yan, H., Zhang, L., Hu, Z., Du, Y., and He,
 538 Y. CODI: Compressing chain-of-thought into contin-
 539 uous space via self-distillation. In Christodoulopoulos,
 540 C., Chakraborty, T., Rose, C., and Peng, V. (eds.), *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pp. 677–693, Suzhou, China, November 2025. Association for Computational Linguistics. ISBN 979-8-89176-332-6. doi: 10.18653/v1/2025.emnlp-main.36. URL <https://aclanthology.org/2025.emnlp-main.36/>.
- 541
 542 Su, D., Zhu, H., Xu, Y., Jiao, J., Tian, Y., and Zheng, Q. To-
 543 ken assorted: Mixing latent and text tokens for improved
 544 language model reasoning. In *Forty-second International Conference on Machine Learning*, 2025. URL <https://openreview.net/forum?id=hYfOPXrbUr>.
- 545
 546 Sui, Y., Chuang, Y.-N., Wang, G., Zhang, J., Zhang, T.,
 547 Yuan, J., Liu, H., Wen, A., Zhong, S., Zou, N., Chen,
 548 H., and Hu, X. Stop overthinking: A survey on efficient
 549 reasoning for large language models, 2025. URL <https://arxiv.org/abs/2503.16419>.
- 550
 551 Tan, W., Li, J., Ju, J., Luo, Z., Song, R., and Luan, J. Think silently, think fast: Dynamic latent compression of LLM reasoning chains. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2025. URL <https://openreview.net/forum?id=AQsko3PPUe>.
- 552
 553 Van Den Oord, A., Vinyals, O., et al. Neural discrete rep-
 554 resentation learning. *Advances in neural information*
 555 *processing systems*, 30, 2017.
- 556
 557 Vincent, P., Larochelle, H., Bengio, Y., and Manzagol, P.-A. Extracting and composing robust features with denoising autoencoders. In *Proceedings of the 25th international conference on Machine learning*, pp. 1096–1103, 2008.
- 558
 559 Wan, F., Huang, X., Cai, D., Quan, X., Bi, W., and Shi,
 560 S. Knowledge fusion of large language models. In *The*

- 550 *Twelfth International Conference on Learning Representations*, 2024. URL [https://openreview.net/](https://openreview.net/forum?id=jiDsk12qcz)
551 [forum?id=jiDsk12qcz](https://openreview.net/forum?id=jiDsk12qcz).
552
553
- 554 Wang, J., Wu, Z., Lai, F., Lian, S., and Zeng, Z. Synadapt:
555 Learning adaptive reasoning in large language models
556 via synthetic continuous chain-of-thought. *arXiv preprint*
557 *arXiv:2508.00574*, 2025.
558
- 559 Wang, Z., Li, X., Xia, R., and Liu, P. Mathpile: A
560 billion-token-scale pretraining corpus for math. In
561 *The Thirty-eight Conference on Neural Information*
562 *Processing Systems Datasets and Benchmarks Track*,
563 2024. URL [https://openreview.net/forum?](https://openreview.net/forum?id=RSvhU69sbG)
564 [id=RSvhU69sbG](https://openreview.net/forum?id=RSvhU69sbG).
565
- 566 Wei, J., Wang, X., Schuurmans, D., Bosma, M., brian ichter,
567 Xia, F., Chi, E. H., Le, Q. V., and Zhou, D. Chain of
568 thought prompting elicits reasoning in large language
569 models. In Oh, A. H., Agarwal, A., Belgrave, D., and Cho,
570 K. (eds.), *Advances in Neural Information Processing*
571 *Systems*, 2022. URL [https://openreview.net/](https://openreview.net/forum?id=_VjQlMeSB_J)
572 [forum?id=_VjQlMeSB_J](https://openreview.net/forum?id=_VjQlMeSB_J).
573
- 574 Yu, J., Li, X., Koh, J. Y., Zhang, H., Pang, R., Qin, J., Ku,
575 A., Xu, Y., Baldrige, J., and Wu, Y. Vector-quantized
576 image modeling with improved vqgan. *arXiv preprint*
577 *arXiv:2110.04627*, 2021.
578
- 579 Yu, L., Jiang, W., Shi, H., Yu, J., Liu, Z., Zhang, Y., Kwok,
580 J. T., Li, Z., Weller, A., and Liu, W. Metamath: Boot-
581 strap your own mathematical questions for large language
582 models. *arXiv preprint arXiv:2309.12284*, 2023a.
583
- 584 Yu, L., Lezama, J., Gundavarapu, N. B., Versari, L., Sohn,
585 K., Minnen, D., Cheng, Y., Birodkar, V., Gupta, A., Gu,
586 X., et al. Language model beats diffusion–tokenizer is key
587 to visual generation. *arXiv preprint arXiv:2310.05737*,
588 2023b.
589
- 590 Yu, Q., Weber, M., Deng, X., Shen, X., Cremers, D., and
591 Chen, L.-C. An image is worth 32 tokens for reconstruc-
592 tion and generation. *Advances in Neural Information*
593 *Processing Systems*, 37:128940–128966, 2024.
594
- 595 Zhang, J., Zhu, Y., Sun, M., Luo, Y., Qiao, S., Du, L.,
596 Zheng, D., Chen, H., and Zhang, N. LightThinker:
597 Thinking step-by-step compression. In Christodoulopou-
598 los, C., Chakraborty, T., Rose, C., and Peng, V.
599 (eds.), *Proceedings of the 2025 Conference on Em-
600 pirical Methods in Natural Language Processing*, pp.
601 13307–13328, Suzhou, China, November 2025. As-
602 sociation for Computational Linguistics. ISBN 979-
603 8-89176-332-6. doi: 10.18653/v1/2025.emnlp-main.
604 673. URL [https://aclanthology.org/2025.](https://aclanthology.org/2025.emnlp-main.673/)
[emnlp-main.673/](https://aclanthology.org/2025.emnlp-main.673/).
- Zhang, R., Isola, P., Efros, A. A., Shechtman, E., and Wang,
O. The unreasonable effectiveness of deep features as a
perceptual metric. In *Proceedings of the IEEE conference*
on computer vision and pattern recognition, pp. 586–595,
2018.
- Zheng, C., Vuong, T.-L., Cai, J., and Phung, D. Movq: Mod-
ulating quantized vectors for high-fidelity image genera-
tion. *Advances in Neural Information Processing Systems*,
35:23412–23425, 2022.

Table 4. The hyper-parameters setting of fine-tuning LLMs.

Initial Learning Rate	Learning Rate Schedule	Optimizer	Batch Size	Training Epochs
1e-5	cosine_with_restarts (warm up steps=all training step * 15%)	AdamW ($\beta_1 = 0.9, \beta_2 = 0.95,$ $weight_decay = 0.1$)	4	5 (50 for GPQA)

Table 5. The value of γ for different LLMs.

Qwen2.5-0.5B-Instruction	Qwen2.5-1.5B-Instruction	LLama3.2-3B-Instruction
-1.58864506	-1.35578818	

A. Experimental Setting

A.1. Benchmarks

We evaluate the proposed method across three representative reasoning scenarios. Mathematical reasoning is first considered and evaluated on GSM8K (Cobbe et al., 2021) and MATH (Hendrycks et al., 2021). GSM8K consists of 8.5K grade-school arithmetic problems, while MATH contains 12.5K challenging competition-level problems. We next evaluate commonsense reasoning on GPQA-Extended (Rein et al., 2024), a scientific question-answering benchmark comprising 546 questions across biology, physics, and chemistry. Finally, we assess logical reasoning using ProsQA (Hao et al., 2024), which contains 17,886 logical reasoning problems. For GSM8K, MATH, and ProsQA, we follow the official training and test splits, while for GPQA-Extended, we adopt the same data split strategy as CoLaR (Tan et al., 2025).

A.2. Implementation

All training and evaluation are conducted on 8 NVIDIA A100 GPUs. Below we provide additional training details beyond those described in the main text.

A.2.1. VISUAL TEXT TOKENIZATION

First, we adopt the text-to-image rendering strategy from Glyph(Cheng et al., 2025). A key difference is that, to fully accommodate an entire CoT within a 512×512 image, we employ a dynamic font-size adjustment strategy. The default font size is set to 9; if the full CoT cannot be rendered at this size, we iteratively reduce the font size until the entire CoT fits within the image. When the font size reaches its minimum and the entire text still cannot be accommodated, we render the remaining content into additional images. Notably, this situation does not arise in the datasets considered in this paper; however, this mechanism can be applied in more extreme scenarios. Conversely, when substantial blank regions are detected, we increase the font size until the blank area occupies no more than 50% of the image.

Then, the pipeline for training VQVAE to conduct visual text tokenization largely follows the official implementation of TiTok (Yu et al., 2024), with both the encoder and decoder consisting of 24 layers transformers. The main differences are that the number of latent tokens n is set to 8, and the embedding dimension d of the codebook is aligned with that of the downstream LLM. Moreover, due to computational constraints, we reduce the batch size in the official implementation to 24. Additional architectural and training details can be found in TiTok.

A.2.2. LLM FINE-TUNING

All relevant training hyperparameters are summarized in Table 4. In addition, Table 5 reports the values of γ of different LLMs. We observe that larger models tend to have higher values of γ , suggesting that γ reflects the model’s level of general knowledge understanding.

B. Case Study for the Ability to Preserve Reasoning Details

First, Figure 6 illustrates the validity of using γ as a threshold to filter unimportant reasoning details. In the first case, the red-highlighted text shows that specialized reasoning skills, such as generating code for diagram construction, are preserved. The second case demonstrates that when no specialized reasoning skills are required, all reasoning steps can be filtered out

```

660 Question: Two boards, one four inches wide and the other six inches wide, are nailed together to form an X. The angle at which they cross is 60 degrees. If this structure is painted and the boards
661 are separated what is the area of the unpainted region on the four-inch board? (The holes caused by the nails are negligible.) Express your answer in simplest radical
662 form.
663 Solution: Note that the unpainted region forms a parallelogram with heights between bases of 4 inches and 6 inches and with one angle 60 degree, as shown.
664 Solution: Note that the unpainted region forms a parallelogram with heights between bases of 4 inches and 6 inches and with one angle 60 degree, as shown.
665 Solution: Note that the unpainted region forms a parallelogram with heights between bases of 4 inches and 6 inches and with one angle 60 degree, as shown.
666 Solution: Note that the unpainted region forms a parallelogram with heights between bases of 4 inches and 6 inches and with one angle 60 degree, as shown.
667 Solution: Note that the unpainted region forms a parallelogram with heights between bases of 4 inches and 6 inches and with one angle 60 degree, as shown.
668 Solution: Note that the unpainted region forms a parallelogram with heights between bases of 4 inches and 6 inches and with one angle 60 degree, as shown.
669 Solution: Note that the unpainted region forms a parallelogram with heights between bases of 4 inches and 6 inches and with one angle 60 degree, as shown.
670 Solution: Note that the unpainted region forms a parallelogram with heights between bases of 4 inches and 6 inches and with one angle 60 degree, as shown.
671 Question: Find all solutions  $SxS$  of the inequality  $S\frac{S}{S}\{24\} + \left\lfloor \frac{S}{48} \right\rfloor < \frac{S}{16}$ . Express your answer in interval notation, simplifying all fractions in your
672 answer.
673 Solution: We can make our work easier by rewriting all fractions in the inequality so that they have a common denominator of 48.  $S\frac{S}{48} - \frac{S}{48} \left\lfloor \frac{S}{48} \right\rfloor < \frac{S}{48}$ . Then we subtract  $\frac{S}{48}$  from both sides:  $S\left\lfloor \frac{S}{48} \right\rfloor < \frac{S}{48} - \frac{S}{48} \left\lfloor \frac{S}{48} \right\rfloor$ . The expression on the left side is the positive difference between  $SxS$  and
674  $\frac{S}{48} \left\lfloor \frac{S}{48} \right\rfloor$ . So, the inequality says that  $SxS$  is strictly between  $\frac{S}{48} \left\lfloor \frac{S}{48} \right\rfloor$  and  $\frac{S}{48} \left\lfloor \frac{S}{48} \right\rfloor + \frac{S}{48}$ . Simplifying these expressions and writing our answer in
675 interval notation, we have  $SxS \in \left( \left\lfloor \frac{S}{48} \right\rfloor, \left\lfloor \frac{S}{48} \right\rfloor + 1 \right)$ .
676 Question: Two distinct integers,  $SxS$  and  $SyS$ , are randomly chosen from the set  $S\{1,2,3,4,5,6,7,8,9,10\}$ . What is the probability that  $Sxy-x-yS$  is even?
677 Solution: We note that  $Sxy-x-yS$  is very close to the expansion of  $S(x-1)(y-1)S$ . (This is basically a use of Simon's Favorite Factoring Trick.) In  $Sxy-x-yS$  is even, then  $Sxy-x-y+1 = (x-1)(y-1)S$  is
678 odd. This only occurs when  $Sx-1S$  and  $Sy-1S$  are both odd, so  $SxS$  and  $SyS$  must be even. There are  $S\binom{5}{2}S$  distinct pairs of even integers, and  $S\binom{10}{2}S$  distinct pairs of integers,
679 so the probability is  $S\frac{\binom{5}{2}}{\binom{10}{2}}S = \frac{1}{9}$ .

```

Figure 6. The three training samples constructed for L-ImgCoT. Strikethrough indicates a reasoning step that is filtered out. The LLM adopted here is Qwen2.5-0.5B-Instruction.

```

681 Question: What is  $S0^{(5^{(6431564)})}S$ ?
682 Solution:  $S0^x=0S$  for any value of  $x$ . ##### 0
683 Question: What is the value of the following expression:  $S100 - 99 + 98 - 97 + 96 - 95 + \dots + 4 - 3 + 2 - 1S$ ?
684 Solution: The given sequence can be rewritten as  $S(100-99) + (98-97) + (96-95) + \dots + (4-3) + (2-1)S$ . ##### 50
685 Question: What is the area of trapezoid SOBCDS below?
686 Solution: We can use the formula for the area of a trapezoid, which is given by  $S\frac{1}{2}(b_1 + b_2)hS$ . To find the height, we draw a line segment from point
687  $SAS$  to the  $x$ -axis, creating a right triangle with legs of length  $SAD = 7S$  and  $SCD = 2S$ . ##### 15

```

Figure 7. The two cases generated by LLM using L-ImgCoT. [-] represents the latent tokens. The LLM adopted here is Qwen2.5-0.5B-Instruction.

under our strategy. The third case shows that our filtering strategy retains specialized skills such as mathematical formula transformations. Together, these cases collectively validate the effectiveness of the proposed filtering strategy.

Figure 7 further illustrates the behaviors exhibited by LLMs trained on the constructed dataset. In the first case, the LLM applies domain-specific knowledge in mathematical reasoning. The second case highlights the LLM’s ability to retain critical details in algebraic formula transformations. In the third case, the LLM explicitly leverages key reasoning details, such as applying the trapezoidal formula and constructing auxiliary lines. Together, these cases demonstrate that LLMs trained with L-ImgCoT are able to preserve critical reasoning details during reasoning.

C. Impact of Latent Token Count

We analyze how model performance varies with an increasing number of latent tokens, as shown in Figure 8. In principle, increasing the number of latent tokens allows more information from the original CoT to be preserved, which is expected to improve reasoning performance. However, the empirical results reveal a non-monotonic trend, where performance first improves and then degrades as the number of latent tokens increases. We hypothesize that, when the number of latent tokens is small, increasing it indeed enriches the preserved CoT information and thus enhances performance. As the number of latent tokens further increases, the combinatorial space of latent representations grows rapidly; with a fixed training dataset, the model tends to underfit and fails to fully capture the underlying semantics. Consistently, when evaluating models with larger numbers of latent tokens, we observe that their latent-token-based reasoning often produces corrupted or incoherent outputs, which provides further evidence supporting our underfitting hypothesis. In future work, we plan to investigate how this trend evolves when the training data scale is increased.

715
716
717
718
719
720
721
722
723
724
725
726
727
728
729
730
731
732
733
734
735
736
737
738
739
740
741
742
743
744
745
746
747
748
749
750
751
752
753
754
755
756
757
758
759
760
761
762
763
764
765
766
767
768
769

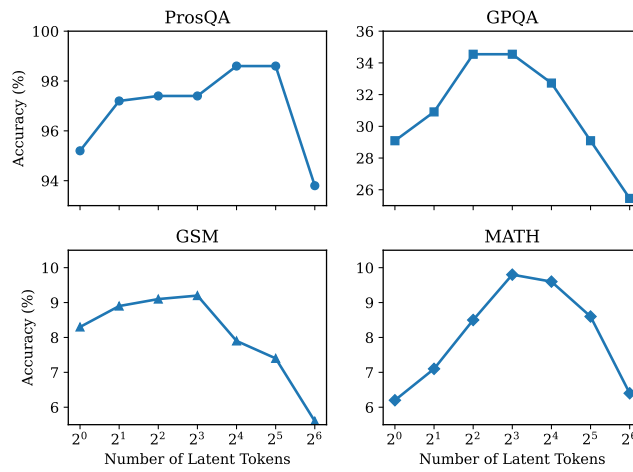


Figure 8. Trend of LLM Reasoning Performance with Varying Numbers of Latent Tokens. The LLM adopted here is Qwen2.5-0.5B-Instruction.