

A STATE-TRANSITION FRAMEWORK FOR EFFICIENT LLM REASONING

Liang Zhang^{1,3}, Yu Zhao², Longyue Wang², Tianqi Shi², Weihua Luo²,
Kaifu Zhang², Jinsong Su^{1,3,4*}

¹School of Informatics, Xiamen University, China ²Alibaba International Digital Commerce

³Key Laboratory of Digital Protection and Intelligent Processing of Intangible Cultural Heritage of Fujian and Taiwan (Xiamen University), Ministry of Culture and Tourism, China

⁴Shanghai Artificial Intelligence Laboratory

lzhang@stu.xmu.edu.cn, jssu@xmu.edu.cn

{fengli.zy, wanglongyue.wly, weihua.luowh}@alibaba-inc.com

ABSTRACT

While Long Chain-of-Thought (CoT) reasoning significantly improves Large Language Models (LLMs) performance on complex reasoning tasks, the substantial computational and memory costs of generating long CoT sequences limit their efficiency and practicality. Existing studies usually enhance the reasoning efficiency of LLMs by compressing CoT sequences. However, this approach conflicts with test-time scaling, limiting the reasoning capacity of LLMs. In this paper, we propose an efficient reasoning framework that models the reasoning process of LLMs as a *state-transition* process. Specifically, we first apply a linear attention mechanism to estimate the LLM’s reasoning state, which records the historical reasoning information from previous reasoning steps. Then, based on the query prompt and the reasoning state, the LLM can efficiently perform the current reasoning step and update the state. With the linear attention, each token in the current reasoning step can directly retrieve relevant historical reasoning information from the reasoning state, without explicitly attending to tokens in previous reasoning steps. In this way, the computational complexity of attention is reduced from quadratic to linear, significantly improving the reasoning efficiency of LLMs. In addition, we propose a state-based reasoning strategy to mitigate the over-thinking issue caused by noisy reasoning steps. Extensive experiments across multiple datasets and model sizes demonstrate that our framework not only improves the reasoning efficiency of LLMs but also enhances their reasoning performance.

1 INTRODUCTION

Chain-of-Thought (CoT) (Wei et al., 2022; Kojima et al., 2022) has become a core technique for enhancing the reasoning ability of large language models (LLMs) on complex tasks. Through prompting step-by-step reasoning, CoT enables LLMs to decompose complex problems into simpler sub-tasks, thus improving their problem-solving capabilities (Yao et al., 2023; Wang et al., 2023; Zhou et al., 2023; Wang et al., 2024). Recent studies, including OpenAI o1 (OpenAI et al., 2024), QwQ (Team, 2024), and DeepSeek-R1 (DeepSeek et al., 2025), demonstrate that scaling up CoT length can further enhance the reasoning abilities of LLMs. However, since most current LLMs are built on the Transformer architecture (Vaswani et al., 2017b), the computational complexity of their attention grows quadratically with context length, and the memory overhead of their KV-cache increases linearly with context length. Hence, generating long CoT substantially increase the computational and memory cost of LLMs, limiting their practical efficiency on complex reasoning tasks.

To improve the reasoning efficiency of LLMs, previous studies employ prompting (Han et al., 2024; Ma et al., 2025a), supervised fine-tuning (SFT) (Liu et al., 2024; Munkhbat et al., 2025), or reinforcement learning (RL) (Aggarwal et al., 2025; Shen et al., 2025) to encourage LLMs toward generating shorter CoT sequences. However, these methods often impair the reasoning ability of

*Corresponding author.

LLMs (Jin et al., 2024; Merrill et al., 2024), since CoT shortening conflicts with test-time scaling (OpenAI et al., 2024). To preserve the reasoning ability of LLMs, some studies (Ma et al., 2025b; Kang et al., 2025; Xia et al., 2025) express the CoT in more concise text (*e.g.*, by removing less important tokens or rewriting with GPT-4) to reduce its length. However, they risk losing critical reasoning information or reducing interpretability when simplifying long CoT (Wang et al., 2025c).

In this paper, we propose an efficient reasoning framework for LLMs, which models the reasoning process of LLMs as a state-transition process. We regard a long CoT as a sequence of reasoning steps, where LLMs perform a specific thinking pattern in each step, such as *induction* or *reflection*. Notably, each reasoning step contains two types of information: substantial linguistic information to ensure its fluency, and limited reasoning information to support subsequent reasoning or answer generation (Zhang et al., 2025a; Xia et al., 2025). Thus, our framework (Figure 1) first compresses the reasoning information from previously generated reasoning steps into a matrix, termed as the *reasoning state matrix*. Then, based on the query prompt and the state matrix, LLMs can efficiently generate the current reasoning step and updates the state matrix accordingly. Specifically, tokens in the current reasoning step can directly retrieve relevant historical reasoning information from the reasoning state, without explicitly attending to tokens in previous steps. In this way, we effectively reduce both the computational complexity of attention and the memory overhead of the KV-cache. Crucially, our framework does not shorten or simplify the CoT sequences generated by LLMs, thus preserving their reasoning ability and interpretability.

To efficiently obtain the state matrix, we design a Mixed Attention Module (**MAM**) to replace the original attention module in LLMs, which consists of a Softmax-Attention (**SA**) submodule and a Linear-Attention (**LA**) submodule. We use the original attention module of LLMs as our SA submodule, where each token can only attend to the tokens in the query prompt and those in its current reasoning step. In the LA submodule, we adopt a linear-attention mechanism (Katharopoulos et al., 2020) to capture the reasoning state of LLMs. Meanwhile, with the linear attention, each token can directly retrieve relevant historical reasoning information from the reasoning state. Compared with other methods, such as CNN (Krizhevsky et al., 2012) and Q-Former (Li et al., 2023), linear attention offers the following advantages in capturing the model’s reasoning state: (1) As a variant of softmax attention, linear attention is naturally compatible with it, thereby reducing the risk of losing critical reasoning information during compression and ensuring the stability and efficiency of model training. (2) Recent studies (Yang et al., 2024b; 2025d) have demonstrated that the state-update process of linear attention is essentially a gradient-descent learning procedure (*i.e.*, test-time training), revealing its strong potential for handling complex reasoning tasks (Zhu et al., 2025).

Another challenge is that LLMs often produce noisy reasoning steps, which may mislead subsequent reasoning steps and lead to the overthinking issue (Chen et al., 2025b; Yang et al., 2025c). To mitigate this issue, we propose a *state-based reasoning strategy*. Given the model’s state-transition process is a gradient-descent process, we first apply the momentum method to accumulate gradients from completed reasoning steps, obtaining a global gradient. The global gradient indicates the global reasoning direction of LLMs. Thus, we employ it to guide LLMs in completing the current reasoning step, ensuring they do not significantly deviate from the global direction.

To validate the efficacy of our framework, we conduct experiments on *seven widely-used benchmark datasets*. Experimental results show that our framework not only improves the reasoning efficiency of LLMs but also enhances their reasoning performance. Extensive ablation studies further demonstrate the effectiveness of various components in our framework.

2 PRELIMINARY

We first present a brief background on linear attention, which lays the foundation for our proposed framework. Meanwhile, we outline the specific form of the CoT sequences in our framework.

2.1 LINEAR ATTENTION

Softmax Attention (SA). Popular LLMs, such as Qwen 3 (Yang et al., 2025a) and Llama 4 (Meta, 2025), adopt a decoder-only Transformer architecture composed of repeated blocks of multi-head softmax attention followed by feed-forward networks (FFNs) (Vaswani et al., 2017a). Given the input sequence $\mathbf{X}=[x_1, \dots, x_{|\mathbf{X}|}]$, the softmax attention can be formulated as follows:

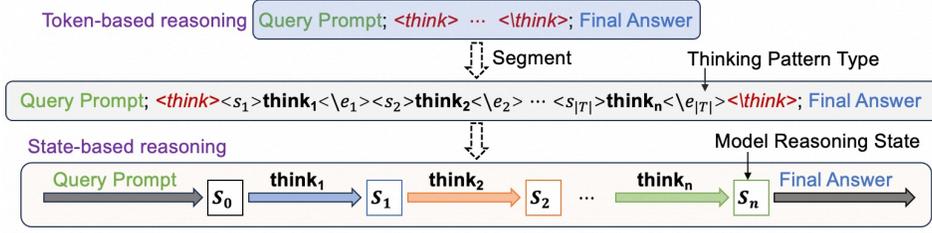


Figure 1: A comparison of traditional *token-based reasoning* with our *state-based reasoning*. \mathbf{think}_t denotes one reasoning step. In state-based reasoning, LLMs can efficiently generate \mathbf{think}_t only based on the query prompt and the reasoning state S_{t-1} .

$$\mathbf{o}_t = \frac{\sum_{i=1}^t \exp(\mathbf{q}_t \mathbf{k}_i^\top / \sqrt{d}) \mathbf{v}_i}{\sum_{i'=1}^t \exp(\mathbf{q}_t \mathbf{k}_{i'}^\top / \sqrt{d})}; \quad \mathbf{q}_t, \mathbf{k}_t, \mathbf{v}_t = x_t \mathbf{W}_Q, x_t \mathbf{W}_K, x_t \mathbf{W}_V, \quad (1)$$

where \mathbf{W}_Q , \mathbf{W}_K , \mathbf{W}_V are learnable weight matrices. The computational complexity of the softmax function increases quadratically with the length of the context sequence. Moreover, softmax attention heavily depends on the growing KV-cache to recall historical information for sequence modeling, leading to substantial memory overheads, particularly in the long context setting.

Linear Attention (LA). Linear attention is a variant of softmax attention, designed to reduce its computational complexity and memory costs. Here, we first replaces the exponential function $\exp(\cdot)$ in softmax attention with a simpler kernel function $\phi(\cdot)$: $\exp(\mathbf{q}_t \mathbf{k}_i / \sqrt{d}) \rightarrow \phi(\mathbf{q}_t) \phi(\mathbf{k}_i)^\top$. Next, based on the associative property of matrix products, linear attention can be written as

$$\mathbf{o}_t = \frac{\sum_{i=1}^t \phi(\mathbf{q}_t) \phi(\mathbf{k}_i)^\top \mathbf{v}_i}{\sum_{i'=1}^t \phi(\mathbf{q}_t) \phi(\mathbf{k}_{i'})^\top} = \frac{\phi(\mathbf{q}_t) \sum_{i=1}^t \phi(\mathbf{k}_i)^\top \mathbf{v}_i}{\phi(\mathbf{q}_t) \sum_{i'=1}^t \phi(\mathbf{k}_{i'})^\top}. \quad (2)$$

Moreover, recent studies (Sun et al., 2023; Yang et al., 2024a) have shown that linear attention can perform well even when $\phi(\cdot)$ is set to the identity function and the normalizer is removed:

$$\mathbf{o}_t = \mathbf{q}_t \sum_{i=1}^t \mathbf{k}_i^\top \mathbf{v}_i = \mathbf{q}_t \mathbf{S}_t; \quad \mathbf{S}_t = \sum_{i=1}^t \mathbf{k}_i^\top \mathbf{v}_i, \quad (3)$$

where \mathbf{S}_t is the state matrix storing historical information. By using \mathbf{S}_t , linear attention can perform sequence modeling in linear time while maintaining constant memory overhead.

The Test-Time Training (TTT) Perspective of LA. Recent works (Sun et al., 2024; Chen et al., 2025b) treat the state matrix \mathbf{S}_t in linear attention as a fast-adapting parameter, updated at every token via lightweight gradient descent. They further introduce an online learning objective for linear attention and provide its closed-form gradient descent solution:

$$\begin{aligned} \text{Objective: } \mathcal{L}(\mathbf{S}) &= -\langle \mathbf{S} \mathbf{k}_t, \mathbf{v}_t \rangle, \\ \text{SGD update: } \mathbf{S}_t &= \mathbf{S}_{t-1} - \beta \nabla \mathcal{L}_t(\mathbf{S}_{t-1}) = \mathbf{S}_{t-1} + \beta \mathbf{v}_t \mathbf{k}_i^\top, \end{aligned} \quad (4)$$

where $\langle \cdot, \cdot \rangle$ denotes the inner product, and β is the learning rate. When β is set to 1, the state update process of linear attention is equivalent to training the state matrix \mathbf{S}_t using the learning objective $\mathcal{L}(\mathbf{S})$. While current studies have yet to provide direct evidence that linear attention can improve the model’s reasoning ability, its theoretical properties suggest significant potential (Zhu et al., 2025).

2.2 LONG CoT SEGMENTATION

A long CoT sample (Q, T, A) usually comprises three parts: a query prompt Q , a thinking sequence T , and a final answer A . Given the query prompt Q , LLMs first perform complex reasoning (T), and then generates the final answer A . During reasoning, LLMs engage in various thinking patterns (e.g., *reflection* and *result verification*) and switch between them using common transitional tokens (e.g., “Wait”, “Hmm”) (Yang et al., 2025c; Wang et al., 2025d). Recent work (Wang et al., 2025b)

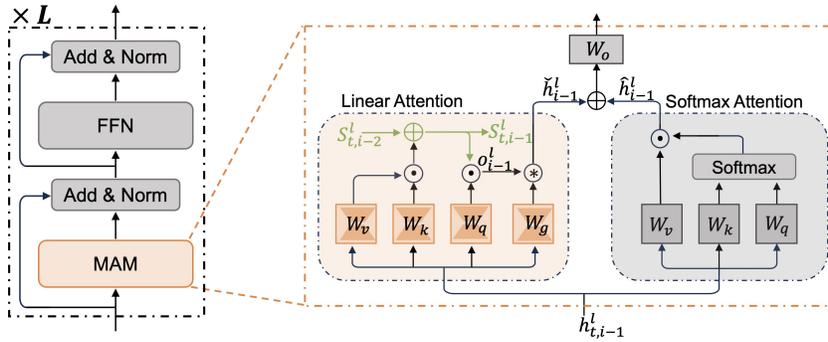


Figure 2: In our framework, the original softmax attention module in LLMs is replaced with our mixed attention module (MAM), thus improving reasoning efficiency and reducing memory cost.

further shows that LLMs usually transition thinking patterns at some high-entropy tokens, such as “*Alternatively*” and “*Maybe*”. Here, we refer to the completion of a thinking pattern by LLMs as a *reasoning step*. Notably, when performing the current reasoning step, LLMs only consider the reasoning information (e.g., conclusion) of previous completed steps without their linguistic information (e.g., grammar).

Following Wang et al. (2025b), we first extract high-entropy transition tokens from the long CoT training set, which occur at the start of a sentence. Next, we use these tokens to segment the thinking sequence T in each training sample into a sequence of reasoning steps. Meanwhile, we cluster all reasoning steps in the entire training set, with each cluster corresponding to a thinking pattern (i.e., type). Finally, for each thinking pattern, we introduce two special tokens to enclose its corresponding reasoning steps, as shown in Figure 1. With the reconstructed training set, the trained LLMs can generate more structured thinking sequences. Meanwhile, these special tokens allow us to track and control the reasoning processes of LLMs more accurately.

3 OUR FRAMEWORK

In this section, we provide a detailed description for our proposed framework. In our framework, we first design a mixed attention module (MAM) to replace the original attention module in LLMs, as illustrated in Figure 2. Using MAM, we can model the reasoning process of LLMs as a state-transition process (see Section 3.1). Then, we further introduces our state-based reasoning strategy in Section 3.2. Finally, our training strategy is presented in Section 3.3.

3.1 MIXED ATTENTION MODULE

Our framework aims to improve the reasoning efficiency of LLMs without sacrificing reasoning ability by modeling their reasoning process as a state-transition process. To achieve this, we design a MAM to *replace* the softmax attention module in LLMs, which consists of a Softmax Attention (SA) submodule and a Linear Attention (LA) submodule. To avoid the performance loss caused by this replacement, we use the original softmax attention module of LLMs as our SA submodule. However, in the SA submodule, each token can only attend to the tokens in the query prompt Q and the previously generated tokens in its reasoning step. By doing so, we reduce the computational complexity of attention from quadratic $O(\mathcal{C}^2)$ to linear $O(\mathcal{C})$ and the memory usage of the KV-cache from linear $O(\mathcal{C})$ to constant $O(1)$, where \mathcal{C} denotes the context length. This significantly improves the reasoning efficiency of our model, especially in complex scenarios requiring longer thinking sequences. Moreover, the LA submodule applies a linear attention mechanism to obtain the LLM’s reasoning state matrix, which records the reasoning information from previously completed reasoning steps. Therefore, each token in the current reasoning step can access relevant historical information from the state matrix without attending directly to tokens in previous reasoning steps.

Given our LLM generates each token in every reasoning step using the same procedure, we take the generation of the i -th token $x_{t,i}$ in the t -th reasoning step as an example. Specifically, at the l -th

layer of our model, we first feed $h_{t,i-1}^{(l-1)}$, the output feature of the input token $x_{t,i-1}$ at the $(l-1)$ -th layer, into the MAM of the current layer. Then, $h_{t,i-1}^{(l-1)}$ is passed to the two submodules in MAM:

In the SA submodule, the KV-cache retains only the KV vectors of the query prompt tokens and ones of the previously generated tokens in the current reasoning step, while those of tokens in completed reasoning steps have been removed. Through the softmax attention mechanism (see Eq. 1), the input token $x_{t,i-1}$ attends to tokens retained in the KV-cache, yielding an updated feature $\hat{h}_{t,i-1}^{(l)}$.

In the LA submodule, after completing the first $t-1$ reasoning steps, our model state is updated to $\mathbf{S}_{t-1}^{(l)}$. Then, we use this state as the initial state $\mathbf{S}_{t,0}^{(l)} = \mathbf{S}_{t-1}^{(l)}$ for the current reasoning step and update it token-by-token, yielding the current model state $\mathbf{S}_{t,i-1}^{(l)} = \mathbf{S}_{t,0}^{(l)} + \sum_{j=1}^{i-1} \mathbf{k}_j^{(l)\top} \mathbf{v}_j^{(l)}$. Next, we utilize the query vector $\mathbf{q}_{i-1}^{(l)}$ of the input token x_{i-1} to extract the relevant historical reasoning information $\mathbf{o}_{i-1}^{(l)}$ from $\mathbf{S}_{t,i-1}^{(l)}$: $\mathbf{o}_{i-1}^{(l)} = \mathbf{q}_{i-1}^{(l)} \mathbf{S}_{t,i-1}^{(l)}$. In the current reasoning step, the model usually relies more on historical reasoning information to generate token in the early stage, and this dependence gradually decreases as more tokens are generated. Therefore, we obtain the output of this submodule via a gating mechanism: $\check{h}_{t,i-1}^{(l)} = \sigma(\mathbf{W}_g h_{t,i-1}^{(l-1)}) * \mathbf{o}_{i-1}^{(l)}$, where $\sigma(\cdot)$ denotes the sigmoid function and \mathbf{W}_g is a learnable weight (see Figure 2).

Finally, we combine the outputs of the two submodules and apply a linear output layer to yield the final output of MAM: $\tilde{h}_{t,i-1}^{(l)} = \mathbf{W}_o(\check{h}_{t,i-1}^{(l)} + \hat{h}_{t,i-1}^{(l)})$. As shown in Figure 2, the two submodules have identical structures, except that the LA submodule incorporates an additional gating weight \mathbf{W}_g . To control its parameter size, we implement the LA submodule via the LoRA strategy (Hu et al., 2022).

Subsequently, we process $\tilde{h}_{t,i-1}^{(l)}$ using the FFN module to produce the output $h_{t,i-1}^{(l)}$ for the l -th layer of our model. By repeating the above process L times, we obtain the final output feature $h_{t,i-1}^{(L)}$ of the input token x_{t-1} , where L denotes the number of layers in our model. Next, $h_{t,i-1}^{(L)}$ is fed into the model’s prediction layer to produce the predicted distribution $p(x)$ of the next token. Finally, our model generates the i -th token x_i of the t -th inference step according to $x_i = \arg \max_x p(x)$.

Following the above procedure, our model sequentially generates tokens in the current reasoning step until reaching its end token $\langle \backslash e_t \rangle$, which corresponds to the thinking pattern of the step (see Figure 1). After generating $\langle \backslash e_t \rangle$, the state matrix of our model at the l -th layer is updated to $\mathbf{S}_{t,n_t}^{(l)}$, which we denote as $\mathbf{S}_t^{(l)}$ and use as the initial state for the next reasoning step, where n_t denotes the number of tokens in the i -th reasoning step. Meanwhile, we clear the KV vectors of all tokens in the current inference step from the KV-cache.

3.2 THE STATE-BASED REASONING STRATEG

During reasoning, LLMs often produce noisy reasoning steps that may mislead subsequent ones, thus resulting in overthinking problems (Chen et al., 2025b; Yang et al., 2025c). In our framework, such noisy reasoning step can deviate the model’s state transitions from the correct reasoning trajectory, resulting in erroneous results (see Figure 3(a)). To mitigate this issue, we propose a state-based reasoning strategy, which guides model reasoning with a global reasoning direction.

In the reasoning process, the state transition of our model at the l -th layer can be formalized as: $\mathbf{S}_0^l \rightarrow \mathbf{S}_1^l \rightarrow \dots \rightarrow \mathbf{S}_{|T|}^l$, where \mathbf{S}_0^l denotes the model state after the LA submodule processes the token in the query prompt \mathbf{Q} . Considering that the state transition process in the line attention is a gradient descent process (see Section 2.1), we represent the total gradient contributed by each reasoning step as $[\nabla_1^l, \nabla_2^l, \dots, \nabla_{|T|}^l]$, where $\nabla_t^l = \mathbf{S}_t^l - \mathbf{S}_{t-1}^l$ indicates the reasoning direction of the t -th step. The reasoning direction of a noisy reasoning step often deviates substantially from those of other steps.

Therefore, we first aggregate the reasoning directions of all previous steps into a global reasoning direction $\bar{\nabla}_{t-1}^l$ using momentum accumulation: $\bar{\nabla}_{t-1}^l = (1 - \frac{1}{t-1}) \bar{\nabla}_{t-2}^l + \frac{1}{t-1} \nabla_{t-1}^l = \frac{1}{t-1} \sum_{i=1}^{t-1} \nabla_i^l$, where $\bar{\nabla}_0^l$ is initialized as a zero matrix. Then, once the t -th reasoning step is completed, we employ the global direction $\bar{\nabla}_{t-1}^l$ to correct its reasoning direction ∇_t^l : $\hat{\nabla}_t^l = (1 - \alpha) \nabla_t^l + \alpha \bar{\nabla}_{t-1}^l$, where $\alpha = \max\{\alpha_{\max}, \frac{t}{|T|}\}$ and $|T|$ denotes the maximum number of reasoning steps (default: 40). Since

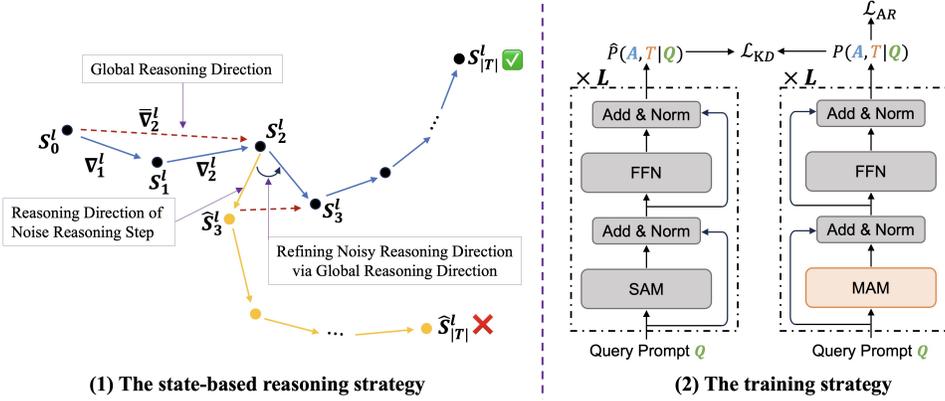


Figure 3: Illustration of our reasoning and training strategies. SAM is the softmax attention module.

more prior reasoning steps yield a more accurate global reasoning direction, we linearly increase the correction coefficient α up to a predefined threshold α_{\max} . In this way, our model can fully explore diverse reasoning directions during the early stages of reasoning and then gradually converge toward the global reasoning direction. Finally, we use the corrected reasoning direction $\hat{\nabla}_t^l$ to update the model state, $S_t^l = S_{t-1}^l + \hat{\nabla}_t^l$, alleviating the negative impact of the noisy step on subsequent steps.

To enhance the diversity of thinking patterns during reasoning, we apply special markers indicating thinking patterns to guide the model toward adopting different thinking patterns across consecutive steps. By doing so, we can further enhance the robustness and overall performance of our model.

3.3 MODEL TRAINING

We train our model using the long CoT training set constructed in Section 2.1. To improve training efficiency while preserving the original reasoning ability of LLMs, we fine-tune only the parameters of the newly added LA submodule and ones of the special tokens corresponding to thinking patterns. As shown in Figure 3(b), we jointly optimize our model with two loss terms: (1) the autoregressive loss \mathcal{L}_{AR} of our model on the training samples, and (2) the knowledge distillation loss \mathcal{L}_{KD} between the base model (the original LLM with softmax attention module) and our proposed model. Finally, our training objective is defined as $\mathcal{L} = \mathcal{L}_{AR} + \beta \mathcal{L}_{KD}$, where β denotes a hyperparameter.

Specifically, we first input a long CoT sample into our model to obtain the predicted probability distribution $P(\mathbf{A}, \mathbf{T} | \mathbf{Q})$ over the thinking sequence \mathbf{T} and the final answer \mathbf{A} conditioned on the query prompt \mathbf{Q} . To maintain the consistency between training and testing, we apply a customized mask matrix in the SA submodule to ensure that each token attends only the tokens in the query prompt \mathbf{A} and those from its corresponding reasoning step. Next, our autoregressive loss term \mathcal{L}_{AR} is defined as follows: $\mathcal{L}_{AR} = -\log P(\mathbf{A}, \mathbf{T} | \mathbf{Q})$.

Meanwhile, the same long CoT sample is fed into the base model to produce the predicted probability distribution $\hat{P}(\mathbf{A}, \mathbf{T} | \mathbf{Q})$. Notably, our model is built upon the base model by replacing its softmax attention module with our MAM. In the base model, each token attends to all previous tokens. Subsequently, we use the Kullback–Leibler (KL) divergence between $P(\mathbf{A}, \mathbf{T} | \mathbf{Q})$ and $\hat{P}(\mathbf{A}, \mathbf{T} | \mathbf{Q})$ as our distillation loss term $\mathcal{L}_{KD} = \text{KL}(\hat{P}(\mathbf{A}, \mathbf{T} | \mathbf{Q}) || P(\mathbf{A}, \mathbf{T} | \mathbf{Q}))$. The \mathcal{L}_{KD} loss term is designed to effectively train our LA submodule for capturing global reasoning information.

4 EXPERIMENTS

4.1 EXPERIMENT SETTINGS

Implementation Details. Our experiments are primarily conducted on the Qwen-2.5 series models (Hui et al., 2024). Meanwhile, we extract 95K high-quality samples from OpenR1-Math-220K to construct our training set using the method described in Section 2.1. We initialize Qwen-2.5 using

Method	GSM8K			MATH-500			AMC23			AIME24			AIME25			AVG.	
	Acc \uparrow	Tok	ReL \downarrow	Acc \uparrow	Tok	ReL \downarrow	Acc \uparrow	Tok	ReL \downarrow	Acc \uparrow	Tok	ReL \downarrow	Acc \uparrow	Tok	ReL \downarrow	Acc \uparrow	ReL \downarrow
<i>Qwen2.5-1.5B Series</i>																	
Ours (Base)	80.1	2086	76.2	78.8	3958	139.4	62.5	6392	242.9	20.0	13765	536.8	16.7	12684	504.7	51.5	300.0
+H2O	76.1	2116	71.7	75.4	3835	128.5	55.0	6230	209.7	13.3	13921	462.2	10.0	12802	428.6	46.0	260.1
+SepLLM	77.3	2230	75.8	76.0	3890	130.3	57.5	6573	220.2	13.3	13690	452.6	13.3	12690	423.9	47.5	260.6
LightThinker	78.8	2109	69.6	77.6	4060	134.8	60.5	6312	214.6	16.7	13827	461.6	13.3	12934	434.6	50.1	263.0
INFITYTHINK	79.5	2503	89.3	78.0	4750	170.2	62.5	7605	263.5	16.7	16318	566.2	16.7	15020	501.0	50.7	318.0
Ours	82.1	2116	69.7	81.2	3812	125.8	67.5	6460	213.2	26.7	13610	440.13	23.3	12910	416.0	56.2	253.0
<i>Qwen2.5-7B Series</i>																	
Ours (Base)	89.4	2649	96.4	87.4	4253	161.6	82.5	6704	242.9	40.0	12901	522.1	30.0	13204	548.1	65.9	314.2
+H2O	83.5	2730	92.0	82.6	4389	147.1	77.5	6843	209.7	30.0	13274	448.9	23.3	13204	443.7	59.4	268.3
+SepLLM	84.9	2582	92.3	84.2	4244	144.3	80.0	6511	220.2	33.3	13100	484.7	23.3	12972	435.9	61.1	275.5
LightThinker	85.6	2779	90.0	84.8	4321	145.2	80.0	6790	214.6	33.3	13307	449.0	26.7	13250	443.9	62.1	268.5
INFITYTHINK	87.9	3237	117.0	86.2	5050	184.8	82.5	8061	254.3	36.7	15094	508.7	30.0	15448	530.5	64.7	319.0
Ours	90.9	2603	87.2	90.0	4196	140.6	85.0	6665	213.2	43.3	13017	434.8	33.3	13191	440.6	68.3	263.3

Table 1: The results of our model and baselines on mathematical reasoning benchmarks.

its corresponding distilled version of DeepSeek-R1 (DeepSeek et al., 2025), and then train it on the training set to obtain **our base model**. Finally, our framework is built upon this base model.

Baselines. We compare our framework with two types of baselines: efficient reasoning methods and KV-cache reduction methods. In the first category, **LightThinker** (Zhang et al., 2025a) uses learnable special tokens to compress the reasoning information of completed reasoning steps, improving the reasoning efficiency of LLMs. **INFITYTHINK** (Yan et al., 2025) compresses the information from previous reasoning steps into a concise summary. In the second category, **H2O** (Zhang et al., 2023) reduces KV-cache by only retaining the KV vectors for a small set of heavy-hitter tokens and recent tokens with a greedy eviction strategy. **SapLLM** (Chen et al., 2025a) only stores the KV vectors of the initial, neighboring, and separator tokens into the KV-cache of LLMs. These baselines are trained with LoRA and use the same number of trainable parameters as our model.

Dataset & Metric. We evaluate our framework on seven benchmarks, including five mathematical reasoning benchmarks (GSM8K (Cobbe et al., 2021), MATH-500 (Hendrycks et al., 2021), AMC23 (AMC, 2023), AIME24 (AIME, 2024), AIME25 (AIME, 2025)), a scientific reasoning benchmark (GPQA_Diamond (Rein et al., 2024)), and a code reasoning benchmark (HumanEval (Chen et al., 2021)). We use greedy decoding to generate the output of the target LLM. We report three metrics to assess the performance and efficiency of various methods: (1) *Accuracy (Acc)* denotes the percentage of correct answers generated by the target model; (2) *Token Number (Tok)* refers to the average length of the CoT sequence; (3) *Reasoning Latency (ReL)* is defined as the average inference time per sample. Moreover, we measure reasoning latency on a single NVIDIA A100 GPU with a batch size of 1, while enabling FlashAttention-2 (Dao, 2023) with BF16 (Kalamkar et al., 2019).

4.2 MAIN RESULTS

The experimental results on mathematical benchmarks are presented in Table 1. As shown in the *AVG.* column, our framework outperforms all baselines in reasoning efficiency and attains the best overall performance. After carefully analyzing these results, we draw several conclusions:

First, our model significantly outperforms the base model in reasoning speed, particularly on benchmarks requiring longer CoT. On the AIME24 benchmark, our model achieved a 16–18% improvement in reasoning speed over the base model. This is mainly attributed to the lower computational complexity of our MAM relative to the softmax attention module of the base model. Furthermore, our model consistently achieves superior performance over the base model across all benchmarks. These results highlight the strong potential of linear attention mechanisms for efficient reasoning.

Second, the baselines (except INFITYTHINK) and our model generate CoT sequences with similar length on each benchmark, as they are trained on the same dataset. While these baselines achieve higher reasoning efficiency than the base model, they exhibit worse reasoning performance. For instance, on the AIME25 benchmark, LightThinker improves reasoning efficiency by 10–12% over the base model, but its performance decreases by 3–6 points. These results suggest that naive KV-cache compression methods risk losing important reasoning information, thus limiting model performance.

Method	GSM8K	MATH-500	AMC23	AIME24	AIME25	AVG.
Ours	82.1	81.2	67.0	26.7	23.3	56.2
(1) <i>w/o</i> The state-based reasoning strategy	79.3	78.8	65.0	20.0	20.0	52.6
(2) <i>w/o</i> Thinking patterns diversity	80.9	80.0	65.0	23.3	20.0	53.8
(3) <i>w/o</i> The distillation loss \mathcal{L}_{KD}	77.0	77.6	62.5	20.0	16.7	50.8
(4) <i>w/o</i> The model reasoning state	76.4	76.2	62.5	20.0	16.7	46.8
(5) <i>w/o</i> The gating mechanism	80.1	79.4	65.0	23.3	23.3	54.2

Table 2: Ablation results of our model on mathematical reasoning benchmarks.

Third, unlike other baselines, INFTYTHINK uses text summaries of previous reasoning steps to record their reasoning information. While enhancing interpretability, this method reduces reasoning efficiency because it requires generating longer CoT sequences. Moreover, such discrete summarization may omit part of the implicit reasoning information within the model, leading to a slight performance drop for INFTYTHINK compared to the base model. However, we effectively avoid these two issues by utilizing the model’s reasoning states to record reasoning information, enabling our model to outperform INFTYTHINK in both efficiency and performance.

4.3 ABLATION STUDY

We further conduct extensive ablation studies by removing different components from our framework to investigate their different impacts. As shown in Table 2, we only compare our model with the following variants in terms of performance, as they share the same reasoning efficiency.

w/o The state-based reasoning strategy. In this variant, we remove the state-based reasoning strategy from our framework. We observe an average performance drop of 3.54 points across all benchmarks for this variant. These results suggest that our reasoning strategy can indeed enhance the robustness of LLMs to noisy reasoning steps, improving their overall performance. Moreover, this also indicates that, with suitable processing, our model states can intuitively reflect the quality of reasoning steps, which offers an interesting insight for future exploration (e.g., process-reward-based RL).

w/o Thinking patterns diversity. We guide LLMs to use different thinking patterns in adjacent reasoning steps via corresponding special tokens, thus ensuring diversity of the thinking patterns in the reasoning process. To test its effectiveness, we remove this operation from our framework. As shown in **Line (2)**, this variant exhibits a consistent decline in performance across all benchmarks. These findings imply that more diverse thinking patterns can enhance LLMs’ exploration ability, thereby improving the accuracy of their generated answers.

The distillation loss \mathcal{L}_{KD} . As described in Section 3.3, we use \mathcal{L}_{KD} to train our LA submodule so that it can better capture global reasoning information. In this variant, \mathcal{L}_{KD} is excluded from our training objective. From **Line (3)** in Table 2, we note that this variant results in a performance drop of 3.6–6.6 points across all benchmarks. The main reason is that the autoregressive loss \mathcal{L}_{AR} may struggle to train the LA submodule to capture the global correlations between reasoning steps.

w/o The model reasoning state. In this variant, we replace the model states in our framework with a sliding window of size 64 to store historical reasoning information. Although this approach is also compatible with the original softmax attention in LLMs, this variant yields the most significant performance drop compared with other variants. This is mainly because this approach lacks reasoning capability as strong as that of linear attention.

w/o The gating mechanism. This variant removes the gating mechanism from our LA submodule and directly sums the outputs of the two submodules in the MAM. This change leads to a slight performance drop across all benchmarks. This confirms that our gating mechanism can more precisely control the amount of historical reasoning information each token requires.

4.4 FURTHER ANALYSIS

Analysis of Computational and Memory Costs. We conduct experiments to further compare the computational and memory efficiency of our model and the base model across varying CoT lengths. The experimental results are presented in Figure 4(a). Although our model exhibits similar reasoning efficiency to the base model for shorter CoT, it significantly surpasses the base model once

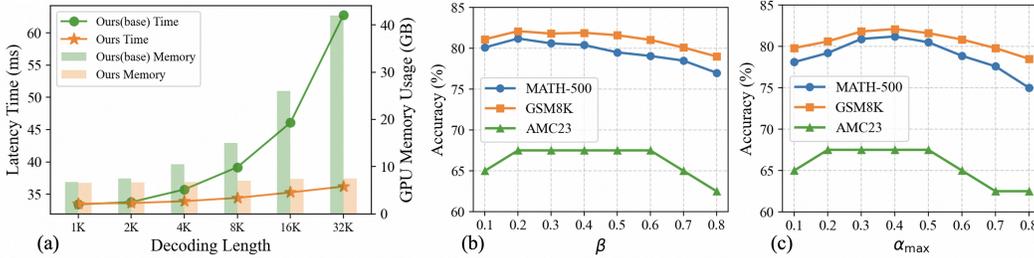


Figure 4: (a) shows the computational and memory efficiency of our model and the base model. (b) and (c) present our model’s performance with different values of hyper-parameters β and α_{\max} , respectively. These experiments are conducted on Qwen2.5-1.5B.

the CoT length exceeds 4K. In particular, when the CoT length reaches 32K, our model achieves over 40% faster reasoning speed than the base model. Moreover, our model maintains a nearly constant memory usage across varying CoT lengths, whereas that of the base model increases linearly with CoT length. Theoretically, our model’s advantages in computational and memory efficiency would become even more significant when FlashAttention-2 is disabled.

Analysis of Hyper-Parameters. We also investigate the impact of the two key hyper-parameters, β and α_{\max} , on the performance of our model. As illustrated in Figure 4(b)–(c), our model exhibits low sensitivity to these two hyper-parameters. Meanwhile, our model attains the best performance when β and α_{\max} are set to 0.2 and 0.4, respectively. We further analyze the choice of these two hyperparameter values as follows:

- β denotes the weight of the distillation loss term \mathcal{L}_{KD} . Setting β to an excessively small value (e.g., 0.1) prevents our linear attention module from effectively capturing global reasoning information, whereas an overly large value (e.g., 0.8) limits the potential of our model to outperform the base model.
- In the state-based reasoning strategy, α determines the magnitude of correction applied to the reasoning direction of noisy reasoning steps. Hence, setting α too low (e.g., 0.1) increases the model’s sensitivity to noisy reasoning steps, whereas setting it too high (e.g., 0.8) constrains the exploration capabilities of LLMs.
- Moreover, within an appropriate range (e.g., 0.2–0.6), our model’s performance is insensitive to these hyperparameters.

Analysis of More Baselines. Recently, many studies (Cheng et al., 2025; Wu et al., 2025; Aggarwal & Welleck, 2025; Li et al., 2025a; Fatemi et al., 2025; Li et al., 2025b; Lan et al., 2025b; Lin et al., 2025; Yuan et al., 2025; Dai et al., 2025a; Xiang et al., 2025; Qi et al., 2025a; Zhang et al., 2025b; Wang et al., 2025a; Yong et al., 2025; Zhao et al., 2025; Chen et al., 2025c; Qi et al., 2025b; Yang et al., 2025b) have focused on shortening the length of CoT to improve the reasoning efficiency of LLMs. However, our method aims to reduce both the computational and storage complexity of attention in LLMs, leading to faster token generation. This indicates that our method improves the reasoning efficiency of LLMs regardless of the CoT length. Therefore, these baselines are both orthogonal to and compatible with our method. Here, we select four open-source baselines as examples to verify the compatibility of our method with these baselines: Task Arithmetic-based Model Merging (TAMM) (Wu et al., 2025), SBT-D (Zhao et al., 2025), L1-Max (Aggarwal & Welleck, 2025), and Anytime-Reasoning (AR) (Qi et al., 2025b). As shown in Table 3, our method further enhances the reasoning efficiency of these baselines by accelerating the token generation speed of LLMs. Furthermore, RL-based baselines (e.g., AR) provide stronger base models for our method, thus further raising its performance upper bound.

In Appendix A, we further analyze both the domain generalization capability of our framework and the gating mechanism within our LA submodule. In Appendix C, we present a more intuitive explanation of our framework and further examine its potential.

Method	GSM8K			MATH-500			AMC23			AIME24			AIME25			AVG.	
	Acc \uparrow	Tok	ReL \downarrow	Acc \uparrow	Tok	ReL \downarrow	Acc \uparrow	Tok	ReL \downarrow	Acc \uparrow	Tok	ReL \downarrow	Acc \uparrow	Tok	ReL \downarrow	Acc \uparrow	ReL \downarrow
<i>Qwen2.5-1.5B Series</i>																	
Ours (Base)	80.1	2086	76.2	78.8	3958	139.4	62.5	6392	242.9	20.0	13765	536.8	16.7	12684	504.7	51.5	300.0
+TAMM	78.0	1212	43.4	73.2	2265	82.1	55.0	2924	105.6	16.7	5702	215.8	16.7	5010	189.0	47.9	127.2
+SBT-D	77.8	1170	43.1	72.8	2105	76.7	57.5	3215	114.9	16.7	6409	246.4	16.7	6173	235.5	48.3	143.3
Ours	82.1	2116	69.7	81.2	3812	125.8	67.5	6460	213.2	26.7	13610	440.13	23.3	12910	416.0	56.2	253.0
+TAMM	82.1	2116	69.7	81.2	3812	125.8	67.5	6460	213.2	26.7	13610	440.13	23.3	12910	416.0	56.2	253.0
+SBT-D	79.9	1081	36.2	76.8	2110	70.3	62.5	3193	101.4	23.3	6458	212.6	20.0	6106	201.5	52.1	124.4
<i>Token Budget</i>	2k			4k			6k			13k			12k				
Base+L1-Max	80.5	2077	77.4	79.0	4091	144.8	62.5	6103	233.5	20.0	13108	514.5	16.7	11904	475.5	51.7	289.1
Base+AR	79.8	2061	75.8	79.4	4083	143.6	62.5	6097	232.0	23.3	13108	514.5	20.0	12088	483.8	53.0	289.9
Ours+L1-Max	82.7	2103	69.5	81.4	3962	129.1	65.0	6025	199.2	26.7	12980	415.2	23.3	11871	382.0	55.8	239.0
Ours+AR	82.7	2084	68.6	82.0	4079	131.2	67.5	6101	200.5	30.0	12980	415.2	26.7	12101	389.3	57.8	240.9
<i>Token Budget</i>	1k			2k			3k			6.5k			6k				
Base+L1-Max	77.8	1161	41.6	76.8	2171	80.3	57.5	3206	124.6	16.7	6606	260.7	16.7	6075	230.8	49.1	147.6
Base+AR	78.8	1105	39.8	78.2	2075	77.8	60.0	3112	124.7	20.0	6587	260.3	16.7	6094	231.6	50.7	146.8
Ours+L1-Max	79.4	1130	36.3	79.2	2066	68.6	62.5	3147	103.1	23.3	6691	208.0	23.3	6030	198.1	53.5	122.8
Ours+AR	81.9	1092	35.1	80.4	2107	69.1	65.0	3107	101.6	26.7	6603	207.4	23.3	6088	198.8	55.5	122.4

Table 3: The results of our model combined with baselines on mathematical reasoning benchmarks.

5 RELATED WORK

Efficient Reasoning. Although long CoT enhances the reasoning ability of LLMs, it introduces significant computational overhead. To alleviate this challenge, existing methods for CoT compression can be roughly divided into three categories. **The first category** leverages prompting (Han et al., 2024; Ma et al., 2025a), supervised fine-tuning (SFT) (Liu et al., 2024; Munkhbat et al., 2025), and reinforcement learning (RL) (Aggarwal et al., 2025; Shen et al., 2025; Dai et al., 2025b) to encourage LLMs to generate shorter CoT sequences. Since such CoT reductions conflict with test-time scaling (OpenAI et al., 2024), these methods often impair the reasoning ability of LLMs (Jin et al., 2024; Merrill et al., 2024). **The second category** compresses information from multiple text tokens into a single hidden vector, converting longer discrete CoT sequences into shorter continuous ones to improve the reasoning efficiency of LLMs (Hao et al., 2024; Su et al., 2025). However, such methods typically require altering the input/output patterns of LLMs, increasing their vulnerability to catastrophic forgetting and limiting their performance. **The third category** expresses the CoT in more concise text to reduce its length while preserving the reasoning ability of LLMs (Ma et al., 2025b; Kang et al., 2025; Xia et al., 2025). Nevertheless, when simplifying long CoT, these methods risk losing critical reasoning information or reducing interpretability (Wang et al., 2025c).

Linear Attention. To address the quadratic computational cost of softmax attention (Zhang et al., 2018; 2022), many linear attention methods have been proposed to improve training and inference efficiency. Early studies investigated kernel functions for linear attention (Kasai et al., 2021; Peng et al., 2021). Subsequently, some studies introduced gating mechanisms into linear attention to more effectively control the information in the model state (Yang et al., 2024a; Qin et al., 2024; Lan et al., 2025a). Recent studies has provided theoretical evidence for the TTT property of linear attention (Yang et al., 2024b; Sun et al., 2024; Liu et al., 2025; Chen et al., 2025b). They focus on exploring online learning objectives for linear attention to enhance its sequence modeling capability.

6 CONCLUSION

In this paper, we propose an efficient reasoning framework that models the reasoning process of LLMs as a state-transition process to enhance their reasoning efficiency. In our framework, we design an MAM to replace the original attention module in LLMs. The MAM uses a linear attention mechanism to capture the model’s reasoning state that stores reasoning information from previous reasoning steps. By utilizing the model state, we can significantly reduce the computational complexity and memory cost of attention, thereby improving the reasoning efficiency of LLMs. Furthermore, we design a state-based reasoning strategy to avoid the overthinking issue caused by noisy reasoning steps. Extensive experiments across multiple benchmarks and model sizes demonstrate the efficiency and robustness of our framework.

ACKNOWLEDGEMENTS

The project was supported by National Key R&D Program of China (No. 2022ZD0160501), Natural Science Foundation of Fujian Province of China (No. 2024J011001), and the Public Technology Service Platform Project of Xiamen (No.3502Z20231043). We also thank the reviewers for their insightful comments.

ETHICS STATEMENT

This work adheres to the ICLR Code of Ethics. In this study, no human subjects or animal experimentation was involved. All datasets used, including OpenR1-Math-220K, GSM8K, MATH-500, AMC23, AIME24, AIME25, GPQA_Diamond, and HumanEval, were sourced in compliance with relevant usage guidelines, ensuring no violation of privacy. We have taken care to avoid any biases or discriminatory outcomes in our research process. No personally identifiable information was used, and no experiments were conducted that could raise privacy or security concerns. We are committed to maintaining transparency and integrity throughout the research process.

REPRODUCIBILITY STATEMENT

We have made every effort to ensure that the results presented in this paper are reproducible. The experimental setup, including training steps, model configurations, and hardware details, is described in detail in the paper. Additionally, we use publicly available datasets, such as OpenR1-Math-220K, GSM8K, MATH-500, and HumanEval, to ensure consistent and reproducible evaluation results. We believe these measures will enable other researchers to reproduce our work and further advance the field.

LLM USAGE

In the preparation of this paper, we use Large Language Models (LLMs) solely to aid in writing and polishing the text, including improving clarity, grammar, and readability. LLMs are not used for generating scientific content, experimental design, analysis, or conclusions. All technical ideas, experiments, and results reported in this paper are entirely the work of the authors.

REFERENCES

- Pranjal Aggarwal and Sean Welleck. L1: Controlling how long a reasoning model thinks with reinforcement learning. In *Proceedings of COLM, 2025*.
- Pranjal Aggarwal, Sean Welleck, Pranjal Aggarwal, and Sean Welleck. L1: Controlling how long a reasoning model thinks with reinforcement learning. *arXiv:2503.04697, 2025*.
- AIME. American invitational mathematics examination (aime) aime 2024-i & ii, 2024. URL https://huggingface.co/datasets/Maxwell-Jia/AIME_2024.
- AIME. American invitational mathematics examination (aime) 2025-i & ii, 2025. URL <https://huggingface.co/datasets/opencompass/AIME2025>.
- AMC. American mathematics competitions, 2023. URL https://artofproblemsolving.com/wiki/index.php/AMC_Problems_and_Solutions.
- Guoxuan Chen, Han Shi, Jiawei Li, Yihang Gao, Xiaozhe Ren, Yimeng Chen, Xin Jiang, Zhenguo Li, Weiyang Liu, and Chao Huang. Sepllm: Accelerate large language models by compressing one segment into one separator. In *Proceedings of ICML, 2025a*.
- Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde De Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, et al. Evaluating large language models trained on code. *arXiv:2107.03374, 2021*.

- Xingyu Chen, Jiahao Xu, Tian Liang, Zhiwei He, Jianhui Pang, Dian Yu, Linfeng Song, Qiuzhi Liu, Mengfei Zhou, Zhuosheng Zhang, et al. Do not think that much for $2+3=?$ on the overthinking of o1-like llms. In *Proceedings of ICML*, 2025b.
- Zigeng Chen, Xinyin Ma, Gongfan Fang, Ruonan Yu, and Xinchao Wang. VeriThinker: Learning to Verify Makes Reasoning Model Efficient. *arXiv:2505.17941*, 2025c.
- Xiaoxue Cheng, Junyi Li, Zhenduo Zhang, Xinyu Tang, Wayne Xin Zhao, Xinyu Kong, and Zhiqiang Zhang. Incentivizing Dual Process Thinking for Efficient Large Language Model Reasoning. *arXiv:2505.16315*, 2025.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, et al. Training verifiers to solve math word problems. *arXiv:2110.14168*, 2021.
- Muzhi Dai, Shixuan Liu, and Qingyi Si. Stable Reinforcement Learning for Efficient Reasoning. *arXiv:2505.18086*, 2025a.
- Muzhi Dai, Chenxu Yang, and Qingyi Si. S-GRPO: Early Exit via Reinforcement Learning in Reasoning Models. *arXiv:2505.07686*, 2025b.
- Tri Dao. Flashattention-2: Faster attention with better parallelism and work partitioning. *arXiv:2307.08691*, 2023.
- DeepSeek, Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, et al. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv:2501.12948*, 2025.
- Mehdi Fatemi, Banafsheh Rafiee, Mingjie Tang, and Kartik Talamadupula. Concise reasoning via reinforcement learning. *arXiv:2504.05185*, 2025.
- Tingxu Han, Zhenting Wang, Chunrong Fang, Shiyu Zhao, Shiqing Ma, and Zhenyu Chen. Token-budget-aware llm reasoning. *arXiv:2412.18547*, 2024.
- Shibo Hao, Sainbayar Sukhbaatar, DiJia Su, Xian Li, Zhiting Hu, Jason Weston, and Yuandong Tian. Training large language models to reason in a continuous latent space. *arXiv:2412.06769*, 2024.
- Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. Measuring mathematical problem solving with the math dataset. In *Proceedings of NIPS*, 2021.
- Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, et al. Lora: Low-rank adaptation of large language models. In *Proceedings of ICLR*, 2022.
- Binyuan Hui, Jian Yang, Zeyu Cui, Jiayi Yang, Dayiheng Liu, Lei Zhang, Tianyu Liu, Jiajun Zhang, Bowen Yu, Keming Lu, et al. Qwen2.5 Technical Report. *arXiv:2409.12186*, 2024.
- Mingyu Jin, Qinkai Yu, Dong Shu, Haiyan Zhao, Wenyue Hua, Yanda Meng, Yongfeng Zhang, and Mengnan Du. The impact of reasoning step length on large language models. In *Proceedings of ACL*, 2024.
- Dhiraj Kalamkar, Dheevatsa Mudigere, Naveen Mellempudi, Dipankar Das, Kunal Banerjee, Sasikanth Avancha, Dharma Teja Vooturi, Nataraj Jammalamadaka, Jianyu Huang, Hector Yuen, et al. A study of BFLOAT16 for deep learning training. *arXiv:1905.12322*, 2019.
- Yu Kang, Xianghui Sun, Liangyu Chen, and Wei Zou. C3ot: Generating shorter chain-of-thought without compromising effectiveness. In *Proceedings of AACL*, 2025.
- Jungo Kasai, Hao Peng, Yizhe Zhang, Dani Yogatama, Gabriel Ilharco, Nikolaos Pappas, Yi Mao, Weizhu Chen, and Noah A Smith. Finetuning pretrained transformers into rnns. In *Proceedings of EMNLP*, 2021.

- Angelos Katharopoulos, Apoorv Vyas, Nikolaos Pappas, and François Fleuret. Transformers are rnns: Fast autoregressive transformers with linear attention. In *Proceedings of ICML*, 2020.
- Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. Large language models are zero-shot reasoners. In *Proceedings of NIPS*, 2022.
- Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Proceedings of NIPS*, 2012.
- Disen Lan, Weigao Sun, Jiayi Hu, Jusen Du, and Yu Cheng. Liger: Linearizing Large Language Models to Gated Recurrent Structures. In *Proceedings of ICML*, 2025a.
- Zhibin Lan, Liqiang Niu, Fandong Meng, Jie Zhou, and Jinsong Su. UME-R1: Exploring Reasoning-Driven Generative Multimodal Embeddings. *arXiv:2511.00405*, 2025b.
- Chen Li, Nazhou Liu, and Kai Yang. Adaptive group policy optimization: Towards stable training and token-efficient reasoning. *arXiv:2503.15952*, 2025a.
- Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *Proceedings of ICML*, 2023.
- Zheng Li, Qingxiu Dong, Jingyuan Ma, Di Zhang, Kai Jia, and Zhifang Sui. Selfbudgeter: Adaptive token allocation for efficient llm reasoning. *arXiv:2505.11274*, 2025b.
- Junhong Lin, Xinyue Zeng, Jie Zhu, Song Wang, Julian Shun, Jun Wu, and Dawei Zhou. Plan and Budget: Effective and Efficient Test-Time Scaling on Large Language Model Reasoning. *arXiv:2505.16122*, 2025.
- Bo Liu, Rui Wang, Lemeng Wu, Yihao Feng, Peter Stone, and Qiang Liu. Longhorn: State space models are amortized online learners. In *Proceedings of ICLR*, 2025.
- Tengxiao Liu, Qipeng Guo, Xiangkun Hu, Cheng Jiayang, Yue Zhang, Xipeng Qiu, and Zheng Zhang. Can language models learn to skip steps? In *Proceedings of NIPS*, 2024.
- Wenjie Ma, Jingxuan He, Charlie Snell, Tyler Griggs, Sewon Min, and Matei Zaharia. Reasoning models can be effective without thinking. *arXiv:2504.09858*, 2025a.
- Xinyin Ma, Guangnian Wan, Runpeng Yu, Gongfan Fang, and Xinchao Wang. Cot-valve: Length-compressible chain-of-thought tuning. *arXiv:2502.09601*, 2025b.
- William Merrill, Ashish Sabharwal, William Merrill, and Ashish Sabharwal. The expressive power of transformers with chain of thought. In *Proceedings of ICLR*, 2024.
- AI Meta. The llama 4 herd: The beginning of a new era of natively multimodal ai innovation. <https://ai.meta.com/blog/llama-4-multimodal-intelligence/>, checked on, 2025.
- Tergel Munkhbat, Namgyu Ho, Seo Hyun Kim, Yongjin Yang, Yujin Kim, and Se-Young Yun. Self-training elicits concise reasoning in large language models. *arXiv:2502.20122*, 2025.
- Aaron OpenAI, Jaech, Adam Kalai, Adam Lerer, Adam Richardson, Ahmed El-Kishky, Aiden Low, Alec Helyar, Aleksander Madry, Alex Beutel, Alex Carney, et al. Openai o1 system card. *arXiv:2412.16720*, 2024.
- Hao Peng, Nikolaos Pappas, Dani Yogatama, Roy Schwartz, Noah A Smith, and Lingpeng Kong. Random feature attention. In *Proceedings of ICLR*, 2021.
- Penghui Qi, Zichen Liu, Tianyu Pang, Chao Du, Wee Sun Lee, and Min Lin. Optimizing anytime reasoning via budget relative policy optimization. *arXiv:2505.13438*, 2025a.
- Penghui Qi, Zichen Liu, Tianyu Pang, Chao Du, Wee Sun Lee, and Min Lin. Optimizing anytime reasoning via budget relative policy optimization. *arXiv:2505.13438*, 2025b.
- Zhen Qin, Songlin Yang, Weixuan Sun, Xuyang Shen, Dong Li, Weigao Sun, and Yiran Zhong. Hgrn2: Gated linear rnns with state expansion. In *Proceedings of COLM*, 2024.

- David Rein, Betty Li Hou, Asa Cooper Stickland, Jackson Petty, Richard Yuanzhe Pang, Julien Dirani, Julian Michael, and Samuel R Bowman. Gpqa: A graduate-level google-proof q&a benchmark. In *Proceedings of COLM*, 2024.
- Yi Shen, Jian Zhang, Jiayun Huang, Shuming Shi, Wenjing Zhang, Jiangze Yan, Ning Wang, Kai Wang, Zhaoxiang Liu, and Shiguo Lian. Dast: Difficulty-adaptive slow-thinking for large reasoning models. *arXiv:2503.04472*, 2025.
- DiJia Su, Hanlin Zhu, Yingchen Xu, Jiantao Jiao, Yuandong Tian, and Qinqing Zheng. Token asorted: Mixing latent and text tokens for improved language model reasoning. *arXiv:2502.03275*, 2025.
- Yu Sun, Xinhao Li, Karan Dalal, Jiarui Xu, Arjun Vikram, Genghan Zhang, Yann Dubois, Xinlei Chen, Xiaolong Wang, Sanmi Koyejo, et al. Learning to (learn at test time): Rnns with expressive hidden states. *arXiv:2407.04620*, 2024.
- Yutao Sun, Li Dong, Shaohan Huang, Shuming Ma, Yuqing Xia, Jilong Xue, Jianyong Wang, and Furu Wei. Retentive network: A successor to transformer for large language models. *arXiv:2307.08621*, 2023.
- Qwen Team. Qwq: Reflect deeply on the boundaries of the unknown. 2024.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Proceedings of NIPS*, 2017a.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Proceedings of NIPS*, 2017b.
- Ante Wang, Linfeng Song, Ye Tian, Baolin Peng, Lifeng Jin, Haitao Mi, Jinsong Su, and Dong Yu. Self-consistency boosts calibration for math reasoning. *arXiv preprint arXiv:2403.09849*, 2024.
- Ante Wang, Linfeng Song, Ye Tian, Baolin Peng, Dian Yu, Haitao Mi, Jinsong Su, and Dong Yu. LiteSearch: Efficient Tree Search with Dynamic Exploration Budget for Math Reasoning. In *Proceedings of AAAI*, 2025a.
- Shenzhi Wang, Le Yu, Chang Gao, Chujie Zheng, Shixuan Liu, Rui Lu, Kai Dang, Xionghui Chen, Jianxin Yang, Zhenru Zhang, et al. Beyond the 80/20 rule: High-entropy minority tokens drive effective reinforcement learning for llm reasoning. *arXiv:2506.01939*, 2025b.
- Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. Self-consistency improves chain of thought reasoning in language models. In *Proceedings of ICLR*, 2023.
- Yibo Wang, Li Shen, Huanjin Yao, Tiansheng Huang, Rui Liu, Naiqiang Tan, Jiaying Huang, Kai Zhang, and Dacheng Tao. R1-Compress: Long Chain-of-Thought Compression via Chunk Compression and Search. *arXiv:2505.16838*, 2025c.
- Yue Wang, Qiuzhi Liu, Jiahao Xu, Tian Liang, Xingyu Chen, Zhiwei He, Linfeng Song, Dian Yu, Juntao Li, Zhuosheng Zhang, et al. Thoughts are all over the place: On the underthinking of o1-like llms. *arXiv:2501.18585*, 2025d.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. In *Proceedings of NIPS*, 2022.
- Han Wu, Yuxuan Yao, Shuqi Liu, Zehua Liu, Xiaojin Fu, Xiongwei Han, Xing Li, Hui-Ling Zhen, Tao Zhong, and Mingxuan Yuan. Unlocking efficient long-to-short llm reasoning with model merging. *arXiv preprint arXiv:2503.20641*, 2025.
- Heming Xia, Chak Tou Leong, Wenjie Wang, Yongqi Li, and Wenjie Li. Tokenskip: Controllable chain-of-thought compression in llms. *arXiv:2502.12067*, 2025.
- Violet Xiang, Chase Blagden, Rafael Rafailov, Nathan Lile, Sang Truong, Chelsea Finn, and Nick Haber. Just Enough Thinking: Efficient Reasoning with Adaptive Length Penalties Reinforcement Learning. *arXiv preprint arXiv:2506.05256*, 2025.

- Yuchen Yan, Yongliang Shen, Yang Liu, Jin Jiang, Mengdi Zhang, Jian Shao, and Yueting Zhuang. Infthythink: Breaking the length limits of long-context reasoning in large language models. *arXiv:2503.06692*, 2025.
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, et al. Qwen3 technical report. *arXiv:2505.09388*, 2025a.
- Chenxu Yang, Qingyi Si, Mz Dai, Dingyu Yao, Mingyu Zheng, Minghui Chen, Zheng Lin, and Weiping Wang. Test-time Prompt Intervention. *arXiv preprint arXiv:2508.02511*, 2025b.
- Chenxu Yang, Qingyi Si, Yongjie Duan, Zheliang Zhu, Chenyu Zhu, Qiaowei Li, Zheng Lin, Li Cao, and Weiping Wang. Dynamic Early Exit in Reasoning Models. *arXiv:2504.15895*, 2025c.
- Songlin Yang, Bailin Wang, Yikang Shen, Rameswar Panda, and Yoon Kim. Gated linear attention transformers with hardware-efficient training. In *Proceedings of ICML*, 2024a.
- Songlin Yang, Bailin Wang, Yu Zhang, Yikang Shen, and Yoon Kim. Parallelizing linear transformers with the delta rule over sequence length. In *Proceedings of NIPS*, 2024b.
- Songlin Yang, Jan Kautz, and Ali Hatamizadeh. Gated delta networks: Improving mamba2 with delta rule. In *Proceedings of ICLR*, 2025d.
- Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Tom Griffiths, Yuan Cao, and Karthik Narasimhan. Tree of thoughts: Deliberate problem solving with large language models. In *Proceedings of NIPS*, 2023.
- Xixian Yong, Xiao Zhou, Yingying Zhang, Jinlin Li, Yefeng Zheng, and Xian Wu. Think or Not? Exploring Thinking Efficiency in Large Reasoning Models via an Information-Theoretic Lens. *arXiv:2505.18237*, 2025.
- Hang Yuan, Bin Yu, Haotian Li, Shijun Yang, Christina Dan Wang, Zhou Yu, Xueyin Xu, Weizhen Qi, and Kai Chen. Not All Tokens Are What You Need In Thinking. *arXiv:2505.17827*, 2025.
- Biao Zhang, Deyi Xiong, and Jinsong Su. Accelerating neural transformer via an average attention network. In *Proceedings of ACL*, 2018.
- Biao Zhang, Deyi Xiong, Yubin Ge, Junfeng Yao, Hao Yue, and Jinsong Su. Aan+: Generalized average attention network for accelerating neural transformer. *Journal of Artificial Intelligence Research*, 2022.
- Jintian Zhang, Yuqi Zhu, Mengshu Sun, Yujie Luo, Shuofei Qiao, Lun Du, Da Zheng, Huajun Chen, and Ningyu Zhang. Lightthinker: Thinking step-by-step compression. *arXiv:2502.15589*, 2025a.
- Xuechen Zhang, Zijian Huang, Chenshun Ni, Ziyang Xiong, Jiasi Chen, and Samet Oymak. Making small language models efficient reasoners: Intervention, supervision, reinforcement. *arXiv:2505.07961*, 2025b.
- Zhenyu Zhang, Ying Sheng, Tianyi Zhou, Tianlong Chen, Lianmin Zheng, Ruisi Cai, Zhao Song, Yuandong Tian, Christopher Ré, Clark Barrett, et al. H2o: Heavy-hitter oracle for efficient generative inference of large language models. In *Proceedings of NIPS*, 2023.
- Haoran Zhao, Yuchen Yan, Yongliang Shen, Haolei Xu, Wenqi Zhang, Kaitao Song, Jian Shao, Weiming Lu, Jun Xiao, and Yueting Zhuang. Let LLMs Break Free from Overthinking via Self-Braking Tuning. *arXiv:2505.14604*, 2025.
- Denny Zhou, Nathanael Schärli, Le Hou, Jason Wei, Nathan Scales, Xuezhi Wang, Dale Schuurmans, Claire Cui, Olivier Bousquet, Quoc Le, et al. Least-to-most prompting enables complex reasoning in large language models. In *Proceedings of ICLR*, 2023.
- Rui-Jie Zhu, Tianhao Peng, Tianhao Cheng, Xingwei Qu, Jinfa Huang, Dawei Zhu, Hao Wang, Kaiwen Xue, Xuanliang Zhang, Yong Shan, et al. A survey on latent reasoning. *arXiv:2507.06203*, 2025.

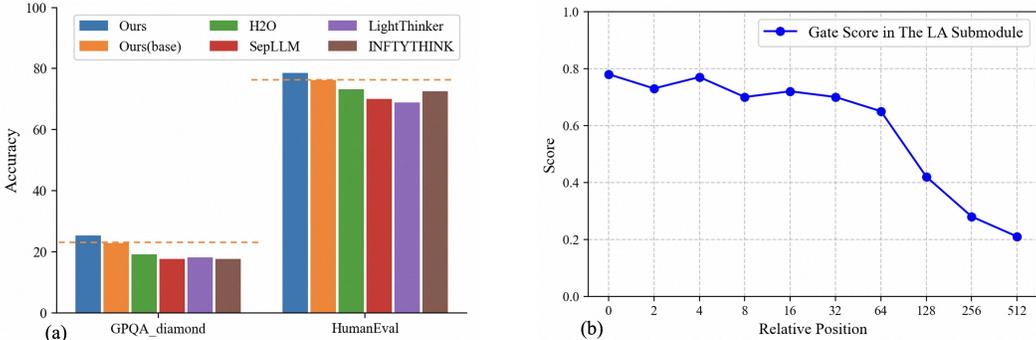


Figure 5: (a) shows the performance of our model and the baselines on the GPQA Diamond and HumanEval benchmarks. (b) illustrates the average gating scores in the LA submodule for tokens at different positions within each reasoning step. These experiments are conducted on Qwen2.5-7B.

A FURTHER ANALYSIS

Analysis of Domain Generalization. As stated in Section 4.1, our training set contains only mathematical reasoning data. To examine the domain generalization of our framework, we further test our model on the scientific reasoning benchmark GPQA_Diamond and the code reasoning benchmark HumanEval. As shown in Figure 5(a), our model consistently outperforms all baselines on these two benchmarks. These results reveal several advantages of our framework: (1) our training strategy effectively prevents the negative impact of target-domain training on the base model’s performance on other domains; (2) our framework can effectively transfer the performance and efficiency benefits in the target domain to other domains.

Analysis of Gating Scores. Here, we further investigate how much historical reasoning information is required by tokens at different positions within each reasoning step. To achieve this, we calculate the average gating scores of these tokens in the LA submodule. As illustrated in Figure 5(b), tokens occurring earlier in a reasoning step generally require more historical reasoning information from the model state than those occurring later. The main reason is that later tokens can directly obtain more information from earlier tokens via the SA sublayer, reducing their reliance on the model state.

B IMPLEMENTATION DETAILS

We train both our model and the baselines for 2 epochs on the training set, with a batch size of 32. The learning rate is set to $2e-4$, with a warmup ratio of 0.03. The learning rate follows a cosine decay schedule to reach zero. We set the learnable LoRA parameters for both our model and the baselines to approximately 66.5M on Qwen2.5-1.5B and 147.3M on Qwen2.5-7B, respectively. To control training cost, we limit the maximum length of long CoT samples in the training set to 16,384 tokens. Meanwhile, we train our model and the baselines on 8 NVIDIA A100 GPUs, each with 80GB of memory. Further implementation details are provided below:

- **First**, to minimize conflicts in thinking types, we employ a K-means clustering algorithm with 128 clusters to annotate the thinking type of each step. A larger number of clusters can effectively reduce the likelihood of assigning reasoning steps with different thinking patterns to the same cluster. Notably, given the large number of reasoning steps (1.4M), we utilize MiniBatch K-Means clustering to improve clustering efficiency. Meanwhile, We use cosine similarity as the distance metric, resulting in a balanced distribution of reasoning steps across clusters (max: 16.1k; min: 4.3k).
- **Second**, we adopt the LoRA strategy to implement our linear attention module, where the rank r is used to control the number of learnable parameters. Specifically, we set the rank r of W_q, W_k, W_v to 512, and the rank r of W_g to 16.

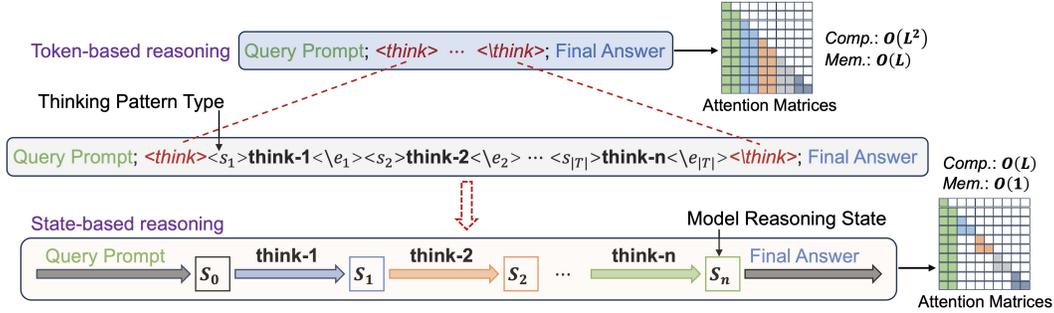


Figure 6: A comparison of traditional *token-based reasoning* with our *state-based reasoning*. \mathbf{think}_t denotes one reasoning step. **Comp.** and **Mem.** represent the computational and memory complexity, respectively, with L indicating the context length. In state-based reasoning, LLMs can efficiently generate \mathbf{think}_t only based on the query prompt and the reasoning state S_{t-1} .

- **Third**, in linear attention, the dimension d of the model state matrix $S_t = S_{t-1} + k_t^T v_t \in R^{d \times d}$ corresponds to the dimensionality of the key and value vectors in LLMs, which is determined by the model architecture itself.
- **Third** as only the newly introduced parameters (the linear attention module and special tokens) are trained, the training cost of our model is reduced to **47.7%** of that of the base model. Concretely, with 8 NVIDIA A100 GPUs and the Qwen-2.5-1.5B model configuration, one training epoch requires 8.6 hours for the base model and 4.1 hours for our model.

C FURTHER ANALYSIS OF OUR FRAMEWORK

In Figure 6, we present a detailed comparison of the reasoning processes between our model (state-based reasoning) and naive LLMs (token-based reasoning). In the case of a naive LLM, we provide it with a query, after which it autoregressively generates the full CoT, where each token attends to all previously generated tokens. Meanwhile, the key and value vectors of all generated tokens must be stored in the KV-cache. Therefore, under this setting, the attention computational complexity scales quadratically with the CoT length L (i.e., $O(L^2)$), while the memory complexity scales linearly (i.e., $O(L)$).

In our framework, LLMs generate the entire CoT step by step, in units of reasoning steps: $\mathbf{think}_1 \rightarrow \dots \rightarrow \mathbf{think}_n$. Meanwhile, we adopt a linear attention mechanism to estimate the state S of LLMs during the reasoning process. As showed in Figure 6 (bottom), as LLM reasoning progresses ($\mathbf{think}_1 \rightarrow \dots \rightarrow \mathbf{think}_n$), its state undergoes continuous transitions ($S_0 \rightarrow \dots \rightarrow S_n$). Therefore, in our framework, the LLM’s reasoning process is modeled as a state-transition process. Importantly, the state of our model performs two key roles:

- **First**, the state records/stores the historical reasoning information of the reasoning steps that have been completed. Hence, our model can generate the next reasoning step (\mathbf{think}_t) based solely on the current state (S_{t-1}) and the query, without depending on any previously completed steps ($\mathbf{think}_{<t}$). In this way, we reduce the computational complexity of attention in LLMs to linear, $O(L)$. Since previously finished reasoning steps are not needed for subsequent reasoning, the key and value vectors of the tokens in these steps are removed from the KV-cache. As a result, the storage complexity of attention in LLMs is reduced to $O(1)$.
- **Second**, the model state can be interpreted as a quantitative representation of the LLM’s knowledge. Consequently, in our state-based reasoning strategy, we correct the reasoning direction of noisy steps to prevent them from contaminating the LLMs’ knowledge (state), thus improving the reasoning performance. Moreover, we are conducting research into using model states to unsupervisedly estimate process rewards for reasoning steps, enhancing the efficiency and effectiveness of LLMs in RL training.

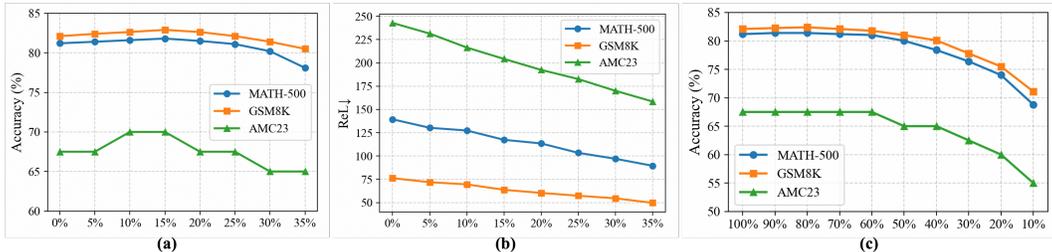


Figure 7: (a) and (b) show the performance and reasoning latency of our model versus the fraction of reasoning steps removed. (c) illustrates our model’s performance after removing varying proportions of training samples. These experiments are conducted on Qwen2.5-1.5B.

Here, we also conducted an analytical experiment to further validate the potential of our model state. Specifically, for each sample in the training set, the process reward of every reasoning step is estimated by computing the cosine of the angle between the global reasoning direction $\bar{\nabla}$ and the step’s reasoning direction ∇_t (i.e., $\cos(\bar{\nabla}, \nabla_t)$). Next, we remove reasoning steps with low process rewards from each sample to shorten the CoT length. Finally, we present the experimental results in Figure 7. We observe that removing less than 20% of the steps resulted in a performance improvement for our model. This indicates that such simple process rewards can reflect the quality of reasoning steps. Moreover, when more than 25% of the steps are removed, our model’s performance begins to decline. This may be attributed to the removal of critical steps in the CoT, which limits the model’s reasoning capacity. Meanwhile, shortening the CoT length further improves the reasoning efficiency of our model.

In addition, we further use our model state to estimate training-sample quality, thereby improving the data efficiency of our method. Specifically, we first calculate the global process reward of a training sample as the average of the process rewards from its reasoning steps. The global process reward reflects the extent to which a sample contains noisy reasoning steps. Then, we discard training samples with low global process rewards and retrain our model. As shown in Figure 7(c), the performance of our model does not degrade when 60% of high-quality training samples are retained. Meanwhile, removing 60% of the training samples (i.e., retaining 40%) results in only a slight performance reduction.

D GENERALIZATION ANALYSIS OF THE CoT SEGMENTATION METHOD

Indeed, our long CoT segmentation method is also a general framework. The framework includes two primary operations:

- **Segmentation**, which aims to segment a long CoT into a sequence of reasoning steps.
- **Annotation**, focused on identifying the thinking type underlying each reasoning step.

Meanwhile, it adheres to two core principles:

- **First**, segmented reasoning steps should maintain low linguistic dependency, mitigating information loss caused by KV-cache cleaning.
- **Second**, conflicts among thinking types should be minimized to avoid detrimental effects from type incompatibilities.

The primary objective of the framework is to enhance the controllability of LLM reasoning processes. In this setting, we can utilize special tokens, associated with the thinking types, to monitor and control the reasoning processes of LLMs.

With sufficient resources, CoT segmentation and thinking-type annotation can be conducted manually, yielding superior performance. An alternative viable approach is to employ powerful LLMs, such as GPT-4o, to carry out segmentation and annotation.

Method	MATH-500			AMC23			AIME24		
<i>Qwen2.5-1.5B Series</i>									
<i>Greedy Decoding</i>	Acc↑	Tok	RC_rate↓	Acc↑	Tok	RC_rate↓	Acc↑	Tok	RC_rate↓
Ours (Base)	78.8	3958	0.31	62.5	6392	0.40	20.0	13765	0.50
Ours	81.2	3812	0.18	67.5	6460	0.23	26.7	13610	0.31
<i>Sampling Decoding</i>	pass@1↑	Tok	RC_rate↓	pass@1↑	Tok	RC_rate↓	pass@1↑	Tok	RC_rate↓
Ours (Base)	80.3	3780	0.19	63.6	6607	0.28	22.5	13880	0.26
Ours	83.5	3652	0.06	69.4	6533	0.08	29.1	13809	0.09

Table 4: The results of our model and baselines under different decoding strategies. RC_rate denotes the proportion of the question for which the model failed to reach the final answer due to repeated generation out of all incorrect questions.

Method	GSM8K			MATH-500			AMC23			AIME24			AIME25			AVG.	
	Acc↑	Tok	ReL↓	Acc↑	Tok	ReL↓	Acc↑	Tok	ReL↓	Acc↑	Tok	ReL↓	Acc↑	Tok	ReL↓	Acc↑	ReL↓
<i>Qwen2.5-7B Series ($\beta = 0.4, \alpha = 0.3$)</i>																	
Ours (Base)	89.4	2649	96.4	87.4	4253	161.6	82.5	6704	242.9	40.0	12901	522.1	30.0	13204	548.1	65.9	314.2
Ours	91.6	2580	86.4	90.4	4244	142.2	87.5	6602	211.7	43.3	12730	432.0	36.7	13310	442.2	69.9	262.9

Table 5: The results of our model and baselines on mathematical reasoning benchmarks.

Given resource limitations, we adopt a more economical strategy: (1) segmentation is performed using high-entropy transition tokens appearing at the beginning of sentences, ensuring that each reasoning step remains semantically independent; and (2) we conduct thinking-type annotation via a clustering algorithm (K-means), while configuring it with a large number of clusters (128) to minimize inter-type conflicts. Nevertheless, our model achieves outstanding performance in the math domain while generalizing effectively to both scientific and code domains.

Furthermore, the details of the clustering algorithm we used are as follows:

- Given the large number of reasoning steps (1.4M), we utilize MiniBatch K-Means clustering to improve clustering efficiency..
- We used a larger number of clusters (128) to effectively reduce the probability of grouping reasoning steps with distinct thinking patterns into the same cluster.
- We use cosine similarity as the distance metric, resulting in a balanced distribution of reasoning steps across clusters (max: 16.1k; min: 4.3k). This effectively reduces the risk of LLMs "over-fitting" or "under-fitting" to thinking patterns.

E FURTHER ANALYSIS OF MODEL PERFORMANCE

Analysis of Decoding Strategies. As observed in DeepSeek-R1 DeepSeek et al. (2025), a temperature-based sampling strategy can effectively reduce the proportion of repeated generation, thus improving the reasoning performance of LLMs. Here, we adopted the DeepSeek-R1 sampling strategy (temperature=0.6, top-k=0.95, run 32 times) and presented pass@1 in Table 4. We observed performance improvements in both our model and the base model. The primary reason is that this sampling method significantly reduces the proportion of repeated generation (i.e., *RC rate*). To avoid overestimating the performance of our model and affecting the reproducibility of our work, we adopted a more deterministic greedy decoding method in the main experiment.

Analysis of Performance-Improvement Magnitude. Notably, the performance gains achieved by our model is non-trivial:

- **First**, in the Qwen-2.5-1.5B scenario, our model achieves an average improvement of 4.7 points over the base model.
- **Second**, although hyperparameters were selected for Qwen-2.5-1.5B, our model delivers an average gain of 2.4 points over the base model in the Qwen-2.5-7B configuration. Nat-

urally, conducting more targeted hyperparameter tuning will further yield greater performance gains (4.0 points) for our model on Qwen-2.5-7B (See Table 5).

- **Third**, the main goal of our work is to improve the reasoning efficiency of LLMs, while performance gains are an unexpected surprise.

F CASE STUDY

In Figure 8 and 9, we present two case studies to illustrate our model’s reasoning process. We note that both cases involve 8 and 9 distinct types of thinking, respectively. Moreover, the two cases also share several thinking types:

- **<type_6>**: problem understanding.
- **<type_25>**: problem interpretation.
- **<type_61>**: algebraic solution.
- **<type_12>**: solution verification.
- ...

We obtain the interpretations for these thinking types using GPT-4o. Moreover, we observe that many reasoning types are general-purpose and not restricted to the mathematical domain, such as **<type_6>**, **<type_25>**, and **<type_12>**. This also indirectly illustrates the domain-generalization capability of our framework.

Question:

Task B-1.3. A ship traveling along a river has covered 24 km upstream and 28 km downstream. For this journey, it took half an hour less than for traveling 30 km upstream and 21 km downstream, or half an hour more than for traveling 15 km upstream and 42 km downstream, assuming that both the ship and the river move uniformly. Determine the speed of the ship in still water and the speed of the river.

Reasoning sequence

<think> **<type_6>** Okay, so I need to find the speed of the ship in still water and the speed of the river. Let me start by recalling that when a ship is moving upstream, its effective speed is the speed of the ship minus the speed of the river. Conversely, when moving downstream, its effective speed is the ship's speed plus the river's speed. **<type_6>** **<type_25>** Let me denote the speed of the ship in still water as v (in km/h) and the speed of the river as r (also in km/h). Then, the upstream speed would be $v - r$, and the downstream speed would be $v + r$. The problem mentions three different journeys: 1. 24 km upstream and 28 km downstream, taking half an hour less than the second journey. 2. 30 km upstream and 21 km downstream, which takes half an hour more than the first journey. 3. 15 km upstream and 42 km downstream, which takes half an hour less than the first journey. Wait, actually, the problem states: "For this journey, it took half an hour less than for traveling 30 km upstream and 21 km downstream, or half an hour more than for traveling 15 km upstream and 42 km downstream..." **<type_25>** **<type_91>** Hmm, let me parse that again. The first journey (24 km upstream, 28 km downstream) took half an hour less than the journey with 30 km upstream and 21 km downstream. Alternatively, it took half an hour more than the journey with 15 km upstream and 42 km downstream. So, there are two comparisons here: Time for 24 up + 28 down = Time for 30 up + 21 down - 0.5 hours; Time for 24 up + 28 down = Time for 15 up + 42 down + 0.5 hours. Therefore, we can set up two equations based on these time differences. First, let me express the time taken for each journey. Time is equal to distance divided by speed. **<type_91>** **<type_78>** So, for the first journey (24 km up, 28 km down): $T_1 = \frac{24}{v-r} + \frac{28}{v+r}$. For the second journey (30 km up, 21 km down): $T_2 = \frac{30}{v-r} + \frac{21}{v+r}$. According to the problem, $T_1 = T_2 - 0.5$. Similarly, for the third journey (15 km up, 42 km down): $T_3 = \frac{15}{v-r} + \frac{42}{v+r}$. And here, $T_1 = T_3 + 0.5$. So, we have two equations: 1. $\frac{24}{v-r} + \frac{28}{v+r} = \frac{30}{v-r} + \frac{21}{v+r} - 0.5$. 2. $\frac{24}{v-r} + \frac{28}{v+r} = \frac{15}{v-r} + \frac{42}{v+r} + 0.5$. Okay, so now I need to solve these two equations for v and r . Let me first work on the first equation: $\frac{24}{v-r} + \frac{28}{v+r} = \frac{30}{v-r} + \frac{21}{v+r} - 0.5$. Let me subtract $\frac{24}{v-r} + \frac{28}{v+r}$ from both sides, but maybe it's better to bring all terms to the left side: $\frac{24}{v-r} + \frac{28}{v+r} - \frac{30}{v-r} - \frac{21}{v+r} + 0.5 = 0$. Simplify the terms: For the upstream terms: $\frac{24-30}{v-r} = \frac{-6}{v-r}$, so $-\frac{6}{v-r}$. For the downstream terms: $\frac{28-21}{v+r} = \frac{7}{v+r}$, so $\frac{7}{v+r}$. So, the equation becomes: $-\frac{6}{v-r} + \frac{7}{v+r} + 0.5 = 0$. Similarly, for the second equation: $\frac{24}{v-r} + \frac{28}{v+r} - \frac{15}{v-r} - \frac{42}{v+r} - 0.5 = 0$. Simplify: Upstream terms: $\frac{24-15}{v-r} = \frac{9}{v-r}$, so $\frac{9}{v-r}$. Downstream terms: $\frac{28-42}{v+r} = \frac{-14}{v+r}$, so $-\frac{14}{v+r}$. Thus: $\frac{9}{v-r} - \frac{14}{v+r} - 0.5 = 0$. **<type_47>** **<type_37>** Now, we have two equations: 1. $-\frac{6}{v-r} + \frac{7}{v+r} + 0.5 = 0$. 2. $\frac{9}{v-r} - \frac{14}{v+r} - 0.5 = 0$. **<type_66>** **<type_66>** Let me denote $x = \frac{1}{v-r}$ and $y = \frac{1}{v+r}$. Then, the equations become: 1. $-6x + 7y + 0.5 = 0$. 2. $9x - 14y - 0.5 = 0$. Now, we have a system of linear equations: Equation (1): $-6x + 7y = -0.5$. Equation (2): $9x - 14y = 0.5$. **<type_47>** **<type_61>** Let me solve this system using the elimination method. Let's multiply Equation (1) by 2 to make the coefficients of y terms opposites: Equation (1) multiplied by 2: $-12x + 14y = -1$. Now, add this to Equation (2): $(-12x + 14y) + (9x - 14y) = -1 + 0.5$. Simplify: $-3x = -0.5$. Thus, $x = \frac{-0.5}{-3} = \frac{0.5}{3} = \frac{1}{6}$. But exact value is $x = 1/6$. Now, substitute $x = 1/6$ back into Equation (1): $-6*(1/6) + 7y = -0.5$. Simplify: $-1 + 7y = -0.5$. Add 1 to both sides: $7y = 0.5$. Thus, $y = 0.5/7 = 1/14$. So, $x = 1/6$ and $y = 1/14$. Recall that $x = \frac{1}{v-r}$ and $y = \frac{1}{v+r}$. $\frac{1}{v-r} = 1/6$ and $\frac{1}{v+r} = 1/14$. Therefore, we have two equations: 1. $v - r = 6$. 2. $v + r = 14$. Adding these equations: $(v - r) + (v + r) = 6 + 14$. Simplify: $2v = 20$. $v = 10$. Substituting back into $v - r = 6$: $10 - r = 6$. $r = 4$. So, the speed of the ship in still water is 10 km/h, and the speed of the river is 4 km/h. **<type_61>** **<type_12>** Let me verify these results by plugging them back into the original time equations. First journey: $24/(10-4) + 28/(10+4) = 24/6 + 28/14 = 4 + 2 = 6$ hours. Second journey: $30/(10-4) + 21/(10+4) = 30/6 + 21/14 = 5 + 1.5 = 6.5$ hours. So, $6.5 - 0.5 = 6$ hours, which matches the first journey. Good. Third journey: $15/(10-4) + 42/(10+4) = 15/6 + 42/14 = 2.5 + 3 = 5.5$ hours. Then, $5.5 + 0.5 = 6$ hours, which also matches the first journey. Perfect. Therefore, the solution is correct. **Final Answer** The speed of the ship in still water is 10 km/h and the speed of the river is 4 km/h. **<type_12>** **</think>**

Figure 8: **Case-1** **<type_i>** and **<type_i>** denote the start and end tokens, respectively, for reasoning steps corresponding to the i -th thinking type.

Question:

3. (6 points) A construction company was building a tunnel. When $\frac{1}{3}$ of the tunnel was completed at the original speed, they started using new equipment, which increased the construction speed by 20% and reduced the working hours to 80% of the original. As a result, it took a total of 185 days to complete the tunnel. If they had not used the new equipment and continued at the original speed, it would have taken \square days to complete the tunnel.

Reasoning sequence:

<think> <type_6> Okay, let me try to figure out this problem. So, a construction company was building a tunnel. When $\frac{1}{3}$ of the tunnel was completed at the original speed, they started using new equipment. This new equipment increased their construction speed by 20% and reduced their working hours to 80% of the original. The total time taken with the new equipment was 185 days. We need to find out how many days it would have taken if they hadn't used the new equipment and continued at the original speed. \n\n <type_6> <type_25> Hmm. Let me break this down. Let's denote the total length of the tunnel as 1 unit for simplicity. So, they completed $\frac{1}{3}$ of the tunnel at the original speed. Then, the remaining $\frac{2}{3}$ was completed with the new equipment, which made them 20% faster and working 80% of the original hours per day. \n\n Wait, the problem says "reduced the working hours to 80% of the original." So, their daily working hours are now 80% of what they used to be. But their speed increased by 20%. So, the rate at which they construct the tunnel per day is affected both by the speed increase and the reduction in working hours. \n\n First, let's define some variables. Let's let the original construction speed be v units per day. <type_25> <type_45> Wait, but actually, maybe it's better to think in terms of time. Let me think. Suppose the original rate is such that without any changes, they would take T days to complete the tunnel. So, total work is 1 tunnel, original time is T days. Therefore, original rate is $\frac{1}{T}$ tunnels per day. \n\n But when they started, they completed $\frac{1}{3}$ of the tunnel at the original speed, then switched to the new equipment for the remaining $\frac{2}{3}$. The total time taken was 185 days. We need to find T . \n\n Wait, but maybe we need to relate the original time to the actual time taken with the changes. Let me approach step by step. \n\n <type_45> <type_70> First, original total time would be T days. But they didn't complete the tunnel in T days; instead, they changed their equipment partway. Let's think about the two phases: Phase 1 is building $\frac{1}{3}$ of the tunnel at original speed, and Phase 2 is building $\frac{2}{3}$ at the new speed with reduced working hours. \n\n Let me define the original rate as R (tunnel per day). Then, the original time to complete the whole tunnel would be $T = \frac{1}{R}$ days. \n\n In Phase 1, they built $\frac{1}{3}$ of the tunnel at rate R . So, the time taken for Phase 1 is $(\frac{1}{3})/R = \frac{1}{3R}$ days. \n\n In Phase 2, they have $\frac{2}{3}$ of the tunnel left. Their speed increases by 20%, so the new speed is $R + 20\%$ of $R = 1.2R$. However, their working hours are reduced to 80% of the original. If the original working hours per day are H , then the new working hours are $0.8H$. But how does this affect the construction rate? \n\n <type_70> <type_78> Wait, speed is in tunnel per day, which is dependent on how much they work each day. If their speed was originally R tunnels per day with H hours per day, then their rate per hour is R/H tunnels per hour. When they increase their speed by 20%, that would be $1.2(R/H)$ tunnels per hour. But they also reduce their working hours to $0.8H$. Therefore, the new daily rate is $1.2(R/H) * 0.8H = 1.2 * 0.8 * R = 0.96R$ tunnels per day. <type_78> <type_101> Wait, is that correct? \n\n Wait, maybe. Let me clarify. The original rate is R tunnels per day, which is based on working H hours per day. So, per hour, they complete R/H tunnels. If their speed (efficiency) increases by 20%, then their per hour rate becomes $1.2(R/H)$. However, they are now working only $0.8H$ hours per day, so the new daily rate is $1.2(R/H) * 0.8H = 1.2 * 0.8 * R = 0.96R$. So, their daily rate actually decreases to $0.96R$? Wait, that seems counterintuitive. If they work faster but fewer hours, the net effect is a decrease in daily rate? <type_101> <type_37> Hmm. \n\n Wait, but 20% faster in speed, but 20% less time working. Let me check the math. If original rate is R per day (H hours). Then per hour rate is R/H . After 20% increase, per hour rate is $1.2R/H$. Then, working $0.8H$ hours per day, so daily rate is $1.2R/H * 0.8H = 0.96R$. So yes, the daily rate becomes $0.96R$. So even though they are working faster, the reduction in hours leads to a lower daily rate. Therefore, the time taken for Phase 2 is $(\frac{2}{3})/0.96R$ days. \n\n But total time is 185 days, so: \n\n Time for Phase 1 + Time for Phase 2 = 185 \n\n Which is: $\frac{1}{3R} + (\frac{2}{3})/(0.96R) = 185$ \n\n We can solve for R , then $T = 1/R$ is the original total time. \n\n <type_37> <type_61> Let me compute that. \n\n First, let's write the equation: $\frac{1}{3R} + (\frac{2}{3})/(0.96R) = 185$ \n\n Combine the terms: $\frac{1}{3} + (\frac{2}{3})/0.96 \bigg/ R = 185$ \n\n Calculate the numerator: $\frac{1}{3} + (\frac{2}{3})/0.96$ \n\n First compute $(\frac{2}{3})/0.96$: $\frac{2}{3} \div 0.96 = (\frac{2}{3}) * (\frac{100}{96}) = (\frac{2}{3}) * (\frac{25}{24}) = (\frac{50}{72}) = \frac{25}{36}$ \n\n Then add $\frac{1}{3}$: $\frac{1}{3} + \frac{25}{36} = \frac{12}{36} + \frac{25}{36} = \frac{37}{36}$ \n\n So total $\frac{37}{36} \div R = 185$ \n\n So numerator $\frac{37}{36}$ \n\n Therefore, $\frac{37}{36} \div R = 185$ \n\n So $R = \frac{37}{36 * 185}$ \n\n So $R = \frac{37}{6660}$ \n\n So $T = 1/R = 185$ days. \n\n Wait, so is T 180 days? Let me check the exact calculation instead of approximating. \n\n Let me compute $\frac{1}{3} + (\frac{2}{3})/0.96$ exactly. \n\n First, $(\frac{2}{3})/0.96$: $\frac{2}{3} \div 0.96$ is $\frac{2}{3} \div \frac{24}{25} = \frac{2}{3} * \frac{25}{24} = \frac{50}{72} = \frac{25}{36}$. \n\n Thus, $\frac{1}{3} + \frac{25}{36} = \frac{12}{36} + \frac{25}{36} = \frac{37}{36}$. \n\n Therefore, numerator is $\frac{37}{36}$. \n\n So the equation is $(\frac{37}{36})/R = 185$ \n\n Hence, $R = \frac{37}{36 * 185} = \frac{37}{6660}$ \n\n $T = 1/R = 185$ days. \n\n Wait, original time would be 180 days. <type_61> <type_12> Wait, so the answer is 180 days? Let me confirm. \n\n If original time is 180 days, then original rate is $\frac{1}{180}$ per day. \n\n Phase 1: $\frac{1}{3}$ of the tunnel at $\frac{1}{180}$ per day. Time = $(\frac{1}{3})/(\frac{1}{180}) = 60$ days. \n\n Phase 2: $\frac{2}{3}$ of the tunnel at the new rate. The new rate is $0.96 * (\frac{1}{180}) = \frac{0.96}{180} = \frac{0.005333...}{1}$ per day. \n\n Time for Phase 2: $(\frac{2}{3})/0.005333... = (\frac{2}{3})/(0.005333...) = \frac{2}{3} * (\frac{1500}{8}) = \frac{2}{3} * 187.5 = 125$ days. \n\n Wait, 60 days + 125 days = 185 days, which matches the given total time. Therefore, the original time T is indeed 180 days. So the answer is 180. \n\n Therefore, if they had not used the new equipment, it would have taken 180 days. \n\n **Final Answer** \n\n \boxed{180} \n\n <type_12> </think>

Figure 9: **Case-2:** <type_i> and <type_i> denote the start and end tokens, respectively, for reasoning steps corresponding to the i -th thinking type.