

Attention Amnesia in Hybrid LLMs: When CoT Fine-Tuning Breaks Long-Range Recall, and How to Fix It

Xinyu Zhou^{1*}, Boyu Zhu^{2*}, Yi Xu³, Zhiwei Li¹, Yingfa Chen⁴,
Huiming Wang^{5†}, Zhijiang Guo^{1,6†}

¹HKUST (GZ), ²UCL, ³Mistral AI, ⁴Tsinghua University, ⁵SUTD, ⁶HKUST

Abstract

Chain-of-thought (CoT) supervised fine-tuning (SFT) is widely adopted to improve reasoning ability, yet we find that it systematically degrades long-context recall in hybrid linear-attention models. Across architectures including HypeNet and Jet-Nemotron, retrieval performance on Needle-In-A-Haystack (NIAH) deteriorates substantially after CoT-SFT, and the degradation becomes more severe under harder retrieval settings and longer context windows. For example, HypeNet-9B on NIAH-S2@256K decreases from 67.2% to 9.4%. We attribute this to CoT-SFT biasing attention gradients toward short-range patterns, disrupting query-key projections (W_Q, W_K) that are responsible for long-range routing. Motivated by this observation, we propose QK-RESTORE, a training-free method that restores only W_Q and W_K from the pre-SFT checkpoint while preserving all other post-SFT parameters. We further introduce a Procrustes variant to balance routing preservation and reasoning adaptation. Across architectures, QK-RESTORE consistently restores long-context capability at zero training cost while preserving reasoning performance; for instance, on HypeNet-5B it improves S3@256K from 65.4% to 76.4% while maintaining strong reasoning performance. Implementation is available at <https://github.com/LARK-AI-Lab/QK-Restore>.

1 Introduction

Efficient sequence models, such as linear attention (Katharopoulos et al., 2020; Qin et al., 2024b; Yang et al., 2024) and state-space models (Gu and Dao, 2024; Dao and Gu, 2024), have emerged as attractive alternatives to softmax-attention (Vaswani et al., 2023) for long-context processing, reducing the quadratic cost to linear complexity through compact recurrent or structured states. However,




Figure 1: NIAH retrieval accuracy before and after CoT-SFT on HypeNet-2B. Despite improving reasoning ability, CoT-SFT substantially degrades long-context retrieval accuracy in distilled hybrid models, especially at longer context lengths and harder NIAH settings.

this compression introduces an information bottleneck on recall-intensive tasks that are crucial for long-context understanding (Zhang et al., 2024), such as Needle-in-a-Haystack (NIAH; Hsieh et al. 2024). Hybrid linear-attention models mitigate this tradeoff by retaining a small set of softmax-attention layers for global recall while converting the remaining layers into efficient linear-attention layers (NVIDIA et al., 2025). Although native hybrid models such as Qwen3.5 (Qwen Team, 2026) and Kimi Linear (Team et al., 2025) achieve strong performance and efficiency, pretraining them from scratch remains highly resource-intensive. This cost motivates recent work that converts pretrained softmax-attention into hybrid linear-attention models through distillation, using only a small fraction of pretraining tokens (Chen et al., 2026; Li et al., 2025b; Hoshino et al., 2025). Existing conversion work primarily focuses on preserving pretraining capabilities during architectural conversion, while the impact of downstream reasoning post-training on long-context retrieval remains underexplored.

While distilled hybrid models primarily serve as efficient base models, modern LLMs typically rely on post-training to acquire stronger instruction-following and reasoning abilities (Guo et al., 2025; Qwen et al., 2025; Ouyang et al., 2022). In particular, CoT-SFT is widely used to enhance multi-

*Equal Contribution.

†Corresponding author.

step reasoning (Wei et al., 2023; Huang et al., 2024; Xu et al., 2025). We observe that applying CoT-SFT on mathematical reasoning data improves reasoning performance, but can substantially degrade the long-context recall ability that distilled hybrid models acquire during pretraining or architectural conversion, particularly in more challenging NIAH settings and extended contexts length (Figure 1). Unlike generic catastrophic forgetting, this degradation is highly structured: it primarily affects long-range retrieval behavior mediated by the retained softmax-attention layers, while leaving the intended reasoning gains largely intact. This reveals a fundamental tension in hybrid models: CoT-SFT strengthens local reasoning can simultaneously disrupt the routing mechanisms required for long-range retrieval.

To investigate this phenomenon, we model CoT reasoning traces as a latent Markov process over intermediate reasoning states, reflecting the local step-to-step structure of mathematical derivations (Wang et al., 2026; Prystawski et al., 2023). Under this CoT-Markov assumption, we derive that the expected gradient magnitude on attention logits decays exponentially with token distance. We empirically validate this prediction by measuring token autocorrelation and gradients of attention-logit. These results suggest that CoT-SFT can improve local multi-step reasoning while eroding the long-range routing behavior needed for retrieval.

To better isolate the source of recall degradation, we analyze the retained softmax-attention layers through a routing-extraction decomposition. We find that CoT-SFT induces locality-biased drift in the query-key projections W_Q, W_K , which determine the source of information retrieval, whereas value-side extraction can still benefit from post-SFT adaptation. Since long-context recall in distilled hybrid models depends heavily on a small number of retained softmax-attention layers, such query-key drift can disproportionately disrupt retrieval. Motivated by this asymmetry, we propose QK-RESTORE, a training-free method that restores only W_Q and W_K in these layers from the pre-SFT checkpoint while preserving all other post-SFT parameters, thereby recovering long-range routing while retaining most reasoning improvements from CoT-SFT. Our contributions can be summarized as:

- We identify CoT-SFT-induced recall degradation as a structured post-training failure mode of distilled hybrid models.

- We provide a theoretical and empirical analysis showing that CoT-SFT concentrates training signals on local token interactions and induces locality-biased drift in the query-key routing geometry of retained softmax-attention layers.
- We introduce QK-RESTORE, a training-free method that restores only query and key projections, recovering long-context recall while largely preserving reasoning performance.

2 Related Work

CoT Dynamics and Attention Routing. Recent work has studied CoT reasoning as a structured generation process over intermediate reasoning states (Wang et al., 2026; Prystawski et al., 2023). A Markovian view formalizes this structure by modeling reasoning as local transitions between latent states, suggesting that CoT supervision primarily reinforces short-range step-to-step dependencies (Wang et al., 2026). In contrast to prior work that utilizes this locality to explain the efficacy of CoT in reasoning, we connect it to a failure mode in efficient long-context models: locality-biased CoT-SFT can erode long-range attention routing required for recall. This perspective aligns with analyses of attention as an information-routing mechanism, where query-key interactions define routing logits that select positions, while value projections determine the content extracted from those positions (Vaswani et al., 2023). We put extended related work in Appendix A.

3 Background and Problem Setup

3.1 Hybrid Model

In this work, we primarily focus on the hybrid model. To retain the expressive power of softmax-attention while improving long-context efficiency, hybrid architectures interleave a small set of softmax-attention layers with efficient linear-attention layers (Yang et al., 2025; Chen et al., 2026).

Softmax-Attention Layers For Transformer’s layer ℓ , head $h \in \{1, \dots, H\}$, the per-head projection matrices $W_Q^{(\ell,h)}, W_K^{(\ell,h)}, W_V^{(\ell,h)} \in \mathbb{R}^{d_h \times d}$ ($d_h = d/H$) transform hidden state $\mathbf{h}_t^{(\ell)} \in \mathbb{R}^d$:

$$\begin{aligned} \mathbf{q}_t^{(\ell,h)} &= W_Q^{(\ell,h)} \mathbf{h}_t^{(\ell)}, & \mathbf{k}_s^{(\ell,h)} &= W_K^{(\ell,h)} \mathbf{h}_s^{(\ell)}, \\ \mathbf{v}_s^{(\ell,h)} &= W_V^{(\ell,h)} \mathbf{h}_s^{(\ell)}, \end{aligned} \quad (1)$$




Figure 2: Overview of QK-RESTORE. After CoT-SFT, the retained softmax-attention layers in a hybrid model may lose long-range routing ability. QK-RESTORE restores only the query and key projections ($\mathbf{W}_Q, \mathbf{W}_K$) from the pre-SFT checkpoint while keeping value-side and other post-SFT parameters unchanged, recovering long-context recall while preserving reasoning performance.

with scalar logit and attention weight

$$e_{ts}^{(\ell,h)} = \frac{(\mathbf{q}_t^{(\ell,h)})^\top \mathbf{k}_s^{(\ell,h)}}{\sqrt{d_h}}, A_{ts}^{(\ell,h)} = \frac{\exp e_{ts}^{(\ell,h)}}{\sum_{u \leq t} \exp e_{tu}^{(\ell,h)}}. \quad (2)$$

The per-head output vector and layer output are

$$\bar{\mathbf{v}}_t = \sum_{s \leq t} A_{ts} \mathbf{v}_s, \mathbf{o}_t^{(\ell)} = \mathbf{W}_O \text{Concat}_h(\bar{\mathbf{v}}_t^{(h)}), \quad (3)$$

Linear-Attention Layers For layer ℓ , a recurrent state matrix $\mathbf{S}_t \in \mathbb{R}^{d_k \times d_v}$ evolves via

$$\mathbf{S}_t^{(\ell)} = \Phi_t^{(\ell)}(\mathbf{S}_{t-1}^{(\ell)}, \mathbf{h}_t^{(\ell-1)}), \mathbf{o}_t^{(\ell)} = \Psi_t^{(\ell)}(\mathbf{S}_t^{(\ell)}, \mathbf{h}_t^{(\ell-1)}), \quad (4)$$

where Φ_t, Ψ_t are layer-specific update and readout operators. Taking Lightning Attention (Qin et al., 2024a) as an example, the state update and readout are:

$$\mathbf{S}_t^{(\ell)} = \text{diag}(\lambda_t^{(\ell)}) \mathbf{S}_{t-1}^{(\ell)} + \mathbf{k}_t^{(\ell)} (\mathbf{v}_t^{(\ell)})^\top, \quad (5)$$

$$\mathbf{o}_t^{(\ell)} = (\mathbf{S}_t^{(\ell)})^\top \mathbf{q}_t^{(\ell)}. \quad (6)$$

We denote the set of softmax-attention layers as L_{attn} . Empirical studies suggest that long-range recall in hybrid models depends disproportionately on a set of softmax-attention layers, whereas most remaining layers can be replaced by recurrent mechanisms with minimal degradation (Wang et al., 2025; Chen et al., 2026; Jelassi et al., 2024). Thus, preserving the routing behavior of L_{attn} is critical for long-context recall.

3.2 Attention Routing

We operationalize attention-routing through the pre-softmax logit $e_{ts} = \mathbf{q}_t^\top \mathbf{k}_s / \sqrt{d_h}$, which determines how strongly position t attends to position

s . Since the routing pattern is determined solely by $\mathbf{W}_Q, \mathbf{W}_K$, changes to these matrices directly alter long-range retrieval behavior. To understand how CoT-SFT training affects retrieval behavior, we examine the gradients w.r.t QK metrics. The gradient update of $\mathbf{W}_Q$ is:

$$\nabla_{\mathbf{W}_Q} \mathcal{L} = \frac{1}{\sqrt{d_h}} \sum_{t,s} \frac{\partial \mathcal{L}}{\partial e_{ts}} \mathbf{k}_s \mathbf{h}_t^\top. \quad (7)$$

If $\partial \mathcal{L} / \partial e_{ts}$ is large for small distance $\tau = t - s$ and negligible for large τ , gradient descent systematically pushes $\mathbf{W}_Q, \mathbf{W}_K$ toward local patterns regardless of context length. Furthermore, we have:

$$\frac{\partial \mathcal{L}}{\partial e_{ts}} = A_{ts} \cdot \mathbf{G}_t^\top (\mathbf{v}_s - \bar{\mathbf{v}}_t), \quad (8)$$

where $\mathbf{G}_t = \partial \mathcal{L} / \partial \mathbf{o}_t \in \mathbb{R}^{d_h}$ is the gradient vector. So understanding the distance dependence of $\partial \mathcal{L} / \partial e_{ts}$ becomes central to understanding how CoT-SFT affects retrieval.

4 Why CoT-SFT Disrupts Routing?

4.1 CoT Data Assumption

We first characterize CoT data structure via a latent Markov model (Wang et al., 2026; Prystawski et al., 2023), which motivates the gradient analysis that follows.

Assumption 4.1 (CoT-Markov structure). There exist latent reasoning states $z_1, \dots, z_T \in \mathcal{Z}$, $|\mathcal{Z}| = K < \infty$, such that:

1. Markov transitions: $P(z_t | z_{t-1}, z_{t-2}, \dots) = P(z_t | z_{t-1})$.




Figure 3: Empirical validation for token autocorrelation decay and attention gradient decay on HypeNet-2B.

2. Observation model: $P(x_t | z_t, x_{t-1}) = P(x_t | z_t)$. The token at position t is determined by the current reasoning state.

3. Ergodicity: The chain is irreducible and aperiodic with stationary distribution π , $\pi_{\min} := \min_{z \in \mathcal{Z}} \pi(z) > 0$.

4. Reversibility: Detailed balance holds: $\pi(z) P(z, z') = \pi(z') P(z', z)$ for all $z, z' \in \mathcal{Z}$.

5. Spectral gap: The transition matrix has second-largest eigenvalue magnitude $\rho \in (0, 1)$.

We include the discussion of the validity of this assumption in [Appendix B](#).

4.2 Gradient Locality Theory

Softmax-attention layers $\ell \in L_{\text{attn}}$ are the sole locus of long-range recall in a hybrid model. Therefore, we analyze how CoT-SFT erodes their routing capacity in these layers through gradient locality.

4.2.1 Definition and Assumptions

Definition 4.2 (Distance-conditioned gradient magnitude).

$$g(\tau) := \mathbb{E}_{t, \ell \in L_{\text{attn}}, h} \left[\left| \frac{\partial \mathcal{L}}{\partial e_{t, t-\tau}^{(\ell, h)}} \right| \right]. \quad (9)$$

Assumption 4.3 (Norm bounds). There exist constants $B_G, B_V > 0$ such that, almost surely over the training distribution,

$$\|G_t\|_2 \leq B_G \quad \text{and} \quad \|v_s\|_2 \leq B_V. \quad (10)$$

Assumption 4.4 (Score-function identity). Let $G_t = \partial \mathcal{L} / \partial o_t$ denote the gradient of the loss with respect to the per-head output at position t . At approximate optimality under cross-entropy, the model satisfies, for all t and almost surely over $x_{1:T}$,

$$\mathbb{E}[G_t | x_{1:t}] = \mathbf{0}. \quad (11)$$

The justifications for the above assumptions are included in [Appendix C](#) and [Appendix D](#).

4.2.2 Gradient Locality Theorem

Lemma 4.5 (Spectral Correlation Decay). *For a stationary reversible ergodic Markov chain with spectral gap $1 - \rho$, and any mean-zero $f, g \in L^2(\pi)$:*

$$|\mathbb{E}[f(z_t) g(z_s)]| \leq \rho^{t-s} \|f\|_{L^2(\pi)} \|g\|_{L^2(\pi)}, \quad (12)$$

Theorem 4.6 (Gradient locality). *Let $B_G, B_V > 0$ be the constants from Assumption 4.3, and let $\rho \in (0, 1)$ be the spectral radius from Assumption 4.1. Under Assumptions 4.3–4.4 and the CoT-Markov model (Assumption 4.1), for all $\tau \geq 0$:*

$$g(\tau) \leq C' \cdot e^{-\tau/W}, \quad (13)$$

where

$$C' := 2B_GB_V > 0, \quad W := \frac{-1}{\log \rho} > 0. \quad (14)$$

The constant C' depends only on the model's norm bounds; the decay rate W is set entirely by the spectral gap of the latent transition matrix. The full proof is included in [Appendix G](#).

4.3 Empirical Validation

Theorem 4.6 shows that CoT training preferentially reinforces nearby token interactions, causing long-range routing signals to decay with distance. The theoretical chain rests on two empirical claims about data and model:

1. Lemma 4.5's hypothesis: the latent Markov chain governing token generation has a spectral gap $1 - \rho$. However, since the latent states z_t are

unobservable, we instead consider the observed token sequence and use the token autocorrelation $\bar{\rho}(\tau)$ as a proxy for ρ . Concretely, $\bar{\rho}(\tau)$ measures how much more likely a token is to reappear at distance τ than by chance, averaged over all token types weighted by their stationary frequency $\pi(v)$. As proved in [Appendix H](#), $\bar{\rho}(\tau) \leq C \cdot \rho^\tau$ under the hidden markov model.

2. Theorem 4.6’s prediction: the expected gradient magnitude $\mathbb{E}[\|\partial\mathcal{L}/\partial e_{ts}\|]$ decays exponentially in τ , with a smaller effective window for CoT data than for general text. We measure $\mathbb{E}[\|\partial\mathcal{L}/\partial e_{ts}\|]$ at each distance τ directly from the model:

$$\frac{\partial\mathcal{L}}{\partial e_{ts}} = A_{ts} \left(\frac{\partial\mathcal{L}}{\partial A_{ts}} - \sum_{s'} A_{ts'} \frac{\partial\mathcal{L}}{\partial A_{ts'}} \right). \quad (15)$$

3. Curve Fitting. Both $\bar{\rho}(\tau)$ and $g(\tau)$ take the form $C \cdot e^{-\tau/W}$, where $W = -1/\log \rho'$ serves as a natural scale of the decay. Therefore, we fit the model $y(\tau) = C \cdot \rho^\tau = C \cdot e^{-\tau/W}$ via log-linear regression. Taking logarithms linearizes the model:

$$\log y(\tau) = \log C - \frac{1}{W}\tau. \quad (16)$$

4. Interpretation. As shown in [Figure 3](#), the left subfigure presents the token autocorrelation decay parameter W_{corr} , a corpus-level statistic that measures the characteristic distance over which text retains structured, non-random self-similarity. Because mathematical reasoning repeatedly references the same symbols, variables, and formula fragments, CoT text sustains high self-similarity over longer distances ($W_{\text{corr}} \approx 172$ tok) than general prose ($W_{\text{corr}} \approx 115$ tok). The right subfigure shows the attention gradient decay parameter W_{grad} , which characterises the training dynamics by measuring the effective reach of the gradient signal $\mathbb{E}[\|\partial\mathcal{L}/\partial e_{ts}\|]$. Both decay curves are well described by exponential fits, consistent with the theorem’s prediction.

When $W_{\text{grad}} < W_{\text{corr}}$, there exists a band of distances $d \in (W_{\text{grad}}, W_{\text{corr}}]$ where the data demands long-range attention but the training gradient no longer reinforces it. The discussions about the attention gradient decay also observed in the pre-training stage are included [Appendix I](#).

KEY TAKEAWAY-1

CoT-SFT improves short-range reasoning ability in hybrid models, at the cost of long-range retrieval.

5 Method

Theorem 4.6 has established that the gradient on the attention logit e_{ts} decays geometrically in distance $\tau = t - s$. Now we investigate which weight metrics this decay propagates into.

5.1 Routing–Extraction Gradient Decoupling

Theorem 5.1 (Routing-Extraction Gradient Decoupling). *Under CoT-SFT, where the token-generation process is modeled as a stationary reversible ergodic Markov chain with a spectral gap $1 - \rho > 0$, gradient updates to the parameters of softmax-attention layer exhibit the asymmetry:*

[Routing] *The per-pair contribution to $\nabla_{\mathbf{W}_Q}$ from position pairs at distance $\tau = t - s$ satisfies:*

$$\mathbb{E} \left[\left\| \frac{\partial\mathcal{L}}{\partial e_{ts}} \mathbf{k}_s \mathbf{h}_t^\top \right\|_F \right] \leq C_R \cdot \rho^\tau$$

for constant $C_R = C' B_K B_h > 0$. The same bound holds for $\nabla_{\mathbf{W}_K}$.

[Extraction] *The gradient on the value vector at any position s satisfies:*

$$\mathbb{E} \left[\left\| \frac{\partial\mathcal{L}}{\partial \mathbf{v}_s} \right\| \right] \geq \delta_A \cdot c_G > 0,$$

where $\delta_A > 0$ satisfies $A_{tt} \geq \delta_A$ for all t and $c_G > 0$ satisfies $\mathbb{E}[\|\mathbf{G}_s\|] \geq c_G$ for all s .

The detailed proof is in [Appendix J](#). Therefore, CoT-SFT affects routing and knowledge extraction differently. $\mathbf{W}_Q, \mathbf{W}_K$ receive meaningful gradient signal only from short-range pairs, while the gradient reaching $\mathbf{W}_V$ is bounded below uniformly over all positions, independent of context length.

5.2 QK-RESTORE

Theorem 5.1 shows that CoT-SFT affects the two functional components of a softmax-attention layer in fundamentally different ways.

The routing parameters $\mathbf{W}_Q$ and $\mathbf{W}_K$, which control where the model attends, receive gradient signal that decays as ρ^τ with distance: they are shaped exclusively by short-range pairs and progressively lose long-range routing capacity. The extraction parameters $\mathbf{W}_V$ and $\mathbf{W}_O$, which control what is retrieved, receive a gradient uniformly bounded below for every position: they accumulate value-processing improvements.

In other words, CoT-SFT simultaneously corrupts $\mathbf{W}_Q, \mathbf{W}_K$ and improves $\mathbf{W}_V, \mathbf{W}_O$. These

effects are segregated into disjoint parameter sets. We propose QK-RESTORE, which eliminates routing corruption and preserves extraction improvement by transplanting $\mathbf{W}_Q, \mathbf{W}_K$ from the pre-SFT checkpoint while retaining post-SFT $\mathbf{W}_V$ and $\mathbf{W}_O$ (Algorithm 1).

Algorithm 1 QK-RESTORE

Input: Hybrid model’s pre-SFT checkpoint θ_{pre} ; Hybrid model’s post-SFT checkpoint θ_{post} ; softmax-attention layers L_{attn}
Output: Repaired model θ_{rep}
Initialized from post-SFT weights
 $\theta_{\text{rep}} \leftarrow \theta_{\text{post}}$
for $\ell \in L_{\text{attn}}$ **do**
 # Replace with pre-SFT weights
 $\mathbf{W}_Q^{\ell, \text{rep}} \leftarrow \mathbf{W}_Q^{\ell, \text{pre}}, \mathbf{W}_K^{\ell, \text{rep}} \leftarrow \mathbf{W}_K^{\ell, \text{pre}}$
end for
return θ_{rep}

6 Experiments

6.1 Setup

Models. HypeNet (Chen et al., 2026) has recently presented strong long-context recall performance, therefore we focus mainly on this model at scales from 2B to 9B. To ensure a comprehensive analysis, we also include Jet-Nemotron-2B¹.

Training. To investigate the influence of CoT-SFT, we primarily focus on the math domain, which is widely studied and has sufficient high-quality datasets (Zhou et al., 2026; Mitra et al., 2024; Yu et al., 2025). We train HypeNet in different scales on our own in both pre-training and SFT stage, and apply SFT to Jet-Nemotron from their pre-trained checkpoints. More details of datasets and training configuration are in Appendix K.

Evaluation. We evaluate the model’s performance on long-context recall and math reasoning tasks. For long-context recall, we report the accuracy on NIAH. To measure math reasoning, we test the model on GSM8K (Cobbe et al., 2021), MATH500 (Lightman et al., 2023), and AIME24 (Zhang and Math-AI, 2024).

6.2 Main Results

The performance comparisons are in Table 1.

CoT-SFT drives long-context degradation. A key observation is that CoT-SFT often weakens long-context retrieval abilities acquired during

pre-training, particularly under more challenging retrieval settings and longer context windows. While simpler tasks such as NIAH-Single-1 remain nearly saturated, substantial degradation emerges on NIAH-Single-2 and NIAH-Single-3.

For example in Table 1, HypeNet-2B on NIAH-Single-3 at 128K drops from 30.4 to 7.60 after SFT, while HypeNet-9B on NIAH-Single-2 at 256K decreases from 67.2 to 9.40. Moreover, the degradation consistently becomes more severe as context length increases, suggesting that SFT disproportionately disrupts the mechanisms responsible for long-range retrieval.

QK-RESTORE recovers long-context capability. QK-RESTORE consistently mitigates the loss of long-context performance incurred during SFT, with the largest gains observed in configurations where the degradation is most severe.

Specifically in Table 1, for HypeNet-2B, NIAH-Single-3 at 256K improves from 10.6 to 30.2 (+19.6), while HypeNet-9B increases from 22.8 to 42.6 (+19.8). Under OpenThoughts-3 CoT-SFT, HypeNet-9B improves from 22.4 to 45.0 on NIAH-Single-2 at 256K and from 40.6 to 61.2 on NIAH-Single-3 at 256K. In several configurations, QK-RESTORE even surpasses the original pre-training baseline. For example, HypeNet-2B on NIAH-Single-3 at 64K improves from 34.8 (pre-train) to 52.8. This suggests that our approach does not simply recover pre-training behavior; rather, it can preserve retrieval structures from pre-training while still benefiting from task-specific representations introduced during SFT.

QK-RESTORE preserves downstream reasoning performance. Beyond retrieval, we evaluate whether restoring routing parameters affects downstream reasoning capabilities. Across benchmarks, QK-RESTORE largely preserves the gains obtained from CoT-SFT while substantially improving long-context retrieval.

On mathematical tasks, performance remains close to the post-SFT model. On HypeNet-5B, QK-RESTORE incurs only minor changes of -0.6 and -0.4 points on MATH500 and GSM8K, respectively, while recovering long-context retrieval performance. Similar trends are observed across other model scales. These results suggest that routing parameters can be selectively restored to recover long-context retrieval without significantly affecting the task-specific knowledge and reasoning capabilities acquired during post-training.

Retrieval difficulty amplifies long-context degra-

¹<https://huggingface.co/collections/jet-ai/jet-nemotron>

Table 1: NIAH and Math performance comparisons across models after CoT-SFT on the Miromind dataset. $\uparrow$ ($\downarrow$) indicates the improvement (degradation) compared to SFT model.

MODEL	METHOD	NIAH-SINGLE-1				NIAH-SINGLE-2				NIAH-SINGLE-3				MATH500		GSM8K		AIME24	
		32K	64K	128K	256K	32K	64K	128K	256K	32K	64K	128K	256K	Avg@16	Maj@16	Avg@16	Maj@16	Avg@16	Maj@16
Jet-Nemotron-2B	Pre-train	100.0	99.8	100.0	100.0	95.6	91.0	61.0	22.2	57.2	48.0	73.6	40.6	37.5	64.0	52.6	84.3	0.80	3.30
	+SFT	100.0	100.0	100.0	98.6	57.8	28.4	21.0	11.6	34.4	33.6	67.2	27.8	49.4	70.6	71.9	88.0	1.50	3.30
	+QK-RESTORE	100.0	100.0	100.0	98.6	64.6	34.0	22.0	7.60	41.8 _(7.4$\uparrow$)	39.0 _(5.4$\uparrow$)	64.4 _(2.8$\downarrow$)	34.0 _(6.2$\uparrow$)	49.0	69.4	71.2	88.6	1.90	6.70
HypeNet-2B	Pre-train	99.2	97.6	97.2	98.4	75.8	94.0	88.0	56.2	48.8	34.8	30.4	23.4	4.70	25.4	2.20	20.0	0.00	0.00
	+SFT	99.8	99.8	99.2	99.6	82.2	81.8	48.4	41.8	47.4	32.6	7.60	10.6	34.4	55.2	40.8	68.1	0.00	0.00
	+QK-RESTORE	99.4	99.6	99.4	99.8	97.6	91.2	83.2	40.8	56.2 _(8.8$\uparrow$)	52.8 _(20.2$\uparrow$)	30.8 _(23.2$\uparrow$)	30.2 _(19.6$\uparrow$)	33.7	56.6	41.2	69.0	0.20	3.30
HypeNet-5B	Pre-train	100.0	100.0	100.0	100.0	99.4	99.8	97.0	93.4	95.0	94.2	87.2	75.2	4.70	33.8	4.10	39.5	0.00	0.00
	+SFT	100.0	100.0	100.0	100.0	98.8	100.0	91.4	83.6	96.6	96.6	90.6	65.4	49.3	67.0	61.3	81.4	0.60	0.00
	+QK-RESTORE	100.0	100.0	100.0	100.0	99.8	100.0	96.8	85.4	97.0 _(0.4$\uparrow$)	96.4 _(0.2$\downarrow$)	90.2 _(0.4$\downarrow$)	76.4 _(11.0$\uparrow$)	48.7	68.4	60.9	81.8	0.40	3.30
HypeNet-9B	Pre-train	100.0	100.0	100.0	100.0	64.2	95.8	72.4	67.2	53.0	95.0	83.8	52.0	17.6	57.6	15.3	62.4	0.80	6.70
	+SFT	100.0	100.0	100.0	100.0	89.0	64.2	17.4	9.40	98.2	90.0	47.8	22.8	39.8	69.5	62.3	87.8	1.20	10.0
	+QK-RESTORE	100.0	100.0	100.0	100.0	92.6	100.0	44.0	19.6	98.6 _(0.4$\uparrow$)	97.0 _(7.0$\uparrow$)	66.4 _(18.6$\uparrow$)	42.6 _(19.8$\uparrow$)	37.6	68.4	59.3	88.4	0.60	3.30

Table 2: NIAH comparisons across models after Instruction-following (non-CoT) SFT on the Tulu-3 dataset.

MODEL	METHOD	NIAH-SINGLE-1				NIAH-SINGLE-2				NIAH-SINGLE-3			
		32K	64K	128K	256K	32K	64K	128K	256K	32K	64K	128K	256K
HypeNet-5B	Pre-train	100.0	100.0	100.0	100.0	99.4	99.8	97.0	93.4	95.0	94.2	87.2	75.2
	+SFT	99.8	100.0	100.0	100.0	99.8	100.0	99.6	93.6	98.6	97.2	85.4	83.8
	+QK-RESTORE	100.0	100.0	100.0	100.0	99.6	100.0	99.8	93.4	97.4	95.2	85.4	85.2
HypeNet-9B	Pre-train	100.0	100.0	100.0	100.0	64.2	95.8	72.4	67.2	53.0	95.0	83.8	52.0
	+SFT	100.0	100.0	100.0	100.0	86.4	97.8	88.2	78.4	96.2	96.8	95.8	86.8
	+QK-RESTORE	100.0	99.8	100.0	100.0	88.2	97.4	92.8	78.6	96.0	97.4	95.2	88.2

Table 3: NIAH comparisons across models after CoT-SFT on the OpenThoughts-3 dataset.

MODEL	METHOD	NIAH-SINGLE-1				NIAH-SINGLE-2				NIAH-SINGLE-3				LCB-V5
		32K	64K	128K	256K	32K	64K	128K	256K	32K	64K	128K	256K	Avg@8
HypeNet-5B	Pre-train	100.0	100.0	100.0	100.0	99.4	99.8	97.0	93.4	95.0	94.2	87.2	75.2	0.27
	+SFT	100.0	100.0	99.8	100.0	99.8	100.0	99.0	89.6	98.8	97.8	94.2	91.4	6.37
	+QK-RESTORE	100.0	99.8	99.8	100.0	99.8	100.0	100.0	92.8	98.6	97.4	93.2	79.4	6.92
HypeNet-9B	Pre-train	100.0	100.0	100.0	100.0	64.2	95.8	72.4	67.2	53.0	95.0	83.8	52.0	1.66
	+SFT	100.0	100.0	100.0	100.0	91.8	100.0	41.8	22.4	96.2	98.6	76.6	40.6	13.11
	+QK-RESTORE	100.0	100.0	100.0	99.8	94.2	99.8	78.2	45.0	95.0	97.4	90.8	61.2	12.91




Figure 4: Attention map from query t to key s of layer33-head8 in pre-SFT and post-SFT on HypeNet-5B. Particularly, panel 4 demonstrates post-SFT assigning higher weights to nearby tokens, which is consistent with our Theorem, and the effectiveness of QK-RESTORE in restoring the long-range attention behavior.

dation. We further observe that retrieval complexity strongly influences robustness. NIAH-Single-1 remains nearly saturated across training stages, whereas substantially larger gaps emerge in Single-2 and Single-3. This trend indicates that difficult retrieval settings place greater demands on precise long-range token interactions and are therefore more sensitive to the degradation.

7 Analysis

7.1 Analysis on Non-CoT SFT

Our theoretical analysis attributes long-context degradation to the localized optimization dynamics induced by CoT supervision. A natural question is whether this phenomenon arises from post-training in general, or whether it is specific to CoT-style reasoning traces. To answer this question, we compare CoT-SFT with conventional instruction-following SFT using the Tulu-3 dataset.

Shown in Table 2, standard instruction tuning

does not exhibit the same degradation pattern observed under CoT-SFT. Across both HypeNet-5B and HypeNet-9B, long-context retrieval performance is largely preserved and frequently improved after Tulu-3 training. For example, HypeNet-5B improves from 75.2 to 83.8 on NIAH-Single-3 at 256K, while HypeNet-9B improves from 52.0 to 86.8. Similar gains are observed on NIAH-Single-2 across multiple context lengths.

Moreover, instruction-following ability improves substantially after Tulu-3 tuning, as measured by IFEval. Despite these improvements, long-context retrieval remains intact. This contrasts sharply with the behavior of CoT-SFT, where retrieval performance deteriorates under identical architectures and context lengths.

These findings suggest that long-context degradation is not an inherent consequence of post-training. Instead, it is closely associated with CoT-style supervision, supporting our hypothesis that the sequential dependency structure of reasoning traces concentrates gradient updates on short-range interactions and disproportionately affects long-range routing behavior.

KEY TAKEAWAY-2

Non-CoT instruction-following SFT preserves or even improves retrieval for hybrid models.

7.2 Analysis Beyond Mathematical Reasoning

Our primary experiments focus on mathematical reasoning, where high-quality CoT supervision is readily available. To evaluate whether the observed retrieval–reasoning trade-off extends beyond mathematics, we conduct additional experiments using OpenThoughts-3 and evaluate downstream coding performance on LiveCodeBench-V5.

In Table 3, the same qualitative behavior persists in the coding domain. CoT-SFT substantially improves coding capability, increasing LiveCodeBench-V5 from 0.27 to 6.37 on HypeNet-5B and from 1.66 to 13.11 on HypeNet-9B. However, these gains are accompanied by noticeable degradation in long-context retrieval, particularly on challenging NIAH settings.

Applying QK-RESTORE recovers a large fraction of the lost retrieval capability while preserving downstream coding performance. For example, on HypeNet-5B, LiveCodeBench-V5 further improves from 6.37 to 6.92, while NIAH-Single-3 at 256K

recovers from 75.2 to 79.4. Similarly, HypeNet-9B improves from 40.6 to 61.2 on NIAH-Single-3 at 256K while maintaining competitive coding performance ($13.11 \rightarrow 12.91$).

These results suggest that the retrieval degradation induced by CoT-SFT is not specific to mathematical reasoning. Instead, it appears to arise from a more general optimization phenomenon associated with long reasoning traces. The effectiveness of QK-RESTORE across both math and code domains further supports our central hypothesis that long-range retrieval is primarily governed by routing parameters.

7.3 Analysis on Attention Routing Recover

Further, we visualize the attention map to characterize how CoT supervised fine-tuning alters the long-range routing behavior of the softmax-attention layers in HypeNet-5B. We conduct the analysis on the models trained on Miromind dataset.

In detail, we compare the attention matrices of a pre-SFT and a post-CoT-SFT checkpoint of HypeNet-5B across all softmax-attention layers. We then identify the most affected layer via a per-head mean attention distance metric:

$$\bar{d}_h = \mathbb{E}_t \left[\sum_{s \leq t} (t-s) A_{h,t,s} \right], \quad (17)$$

where h indexes the attention head and $A_{h,t,s}$ is the post-softmax weight from query t to key s . Ranking layers by the number of heads whose $\bar{d}_h$ decreases after SFT, Layer 33 emerges as the most affected, with 25 of 32 heads shifting toward shorter attention distances. We therefore focus our analysis on this layer.

We visualize the top recall heads on a synthetic NIAH prompt of 8,192 tokens, with the target fact embedded at 50% context depth. For each head, we display a four-panel figure: the pre-SFT attention matrix $\tilde{A}_{\text{pre}}$ (blue), the post-SFT matrix $\tilde{A}_{\text{post}}$ (red), the QK-RESTORE matrix $\tilde{A}_{\text{QK-RESTORE}}$ (green), and the element-wise log-ratio $\log_2(\tilde{A}_{\text{QK-RESTORE}}/\tilde{A}_{\text{post}})$ (Panel 4), where positive values indicate positions where QK-RESTORE recovers more long-range attention than post-SFT.

Analysis. Figure 4 shows the pattern of Layer-33 Head-8 of HypeNet-5B. From panel 4, we observe that the off-diagonal lower triangle is predominantly red, confirming that QK-RESTORE partially recovers the long-range routing suppressed by




Figure 5: Routing geometry drift in HypeNet-2B.




Figure 6: Routing geometry drift in Jet-Nemotron-2B.

CoT fine-tuning: for most query positions t , QK-RESTORE attends more strongly to distant keys $s \ll t$ than post-SFT does. Second, the lower-right block, where both the query and key positions fall after the needle s^* , shows a blue band along the main diagonal, meaning that post-SFT assigns higher weight to nearby tokens in this region, which matches our hypothesis. These observations demonstrate the effectiveness of QK-RESTORE in mitigating routing collapse and restoring the long-range attention behavior characteristic.

7.4 Ablation on QK-RESTORE

QK-RESTORE transplanting the full $\mathbf{W}_Q, \mathbf{W}_K$ matrices from the pre-SFT checkpoint into the post-SFT model, recovers long-range recall substantially but incurs a math performance drop. The post-SFT $\mathbf{W}_Q$ can be decomposed as:

$$\mathbf{W}_Q^{\text{post}} = \mathbf{W}_Q^{\text{pre}} + \delta \mathbf{W}_Q^{\text{route}} + \delta \mathbf{W}_Q^{\text{math}},$$

where $\delta \mathbf{W}_{\text{route}}$ encodes harmful locality drift and $\delta \mathbf{W}_{\text{math}}$ encodes beneficial math ability adaptation. Since both components are entangled within $\mathbf{W}_Q$, QK-RESTORE discards them together, recovering routing at the cost of erasing the math benefit. As expected, Figure 5 and Figure 6 show a non-negligible drift in $\mathbf{R} = \mathbf{W}_Q \mathbf{W}_K^\top$, where $\Delta \mathbf{R} = \mathbf{R}^{\text{post}} - \mathbf{R}^{\text{pre}}$.

Table 4: Ablation of QK-RESTORE in HypeNet-2B.

METHOD	NIAH-SINGLE-2		NIAH-SINGLE-3		MATH500 Avg@16
	64K	128K	64K	128K	
QK-RESTORE	91.2	83.2	52.8	30.8	33.7
QK-PRO	94.0	79.0	51.8	26.0	34.2
Q-RESTORE	78.4	79.0	50.8	24.4	34.2
K-RESTORE	90.4	71.6	46.2	13.2	34.2
V-RESTORE	57.2	36.6	24.0	15.2	32.1
QK-FROZEN	90.6	76.4	49.4	23.6	33.1

Therefore, we model it as a Procrustes problem:

$$\begin{aligned} \min_{\mathbf{W}_Q^{\text{new}}, \mathbf{W}_K^{\text{new}}} & \left\| \mathbf{W}_Q^{\text{new}} - \mathbf{W}_Q^{\text{post}} \right\|_F^2 + \left\| \mathbf{W}_K^{\text{new}} - \mathbf{W}_K^{\text{post}} \right\|_F^2 \\ \text{s.t. } & \mathbf{W}_Q^{\text{new}} \mathbf{W}_K^{\text{new}\top} = \mathbf{R}_{\text{pre}}, \end{aligned} \quad (18)$$

and we call it as QK-PRO variant (more details are in Appendix O).

Furthermore, we evaluate the roles of each components in attention (i.e., $\mathbf{W}_K, \mathbf{W}_Q, \mathbf{W}_V$), and CoT-SFT with $\mathbf{W}_Q^{\text{pre}}, \mathbf{W}_K^{\text{pre}}$ frozen as a preventive counterpart (denoted as QK-FROZEN), to answer a question: is it preferable to prevent routing drift during training, or to correct it post-hoc?

Analysis. Table 4 validates three key claims. V-RESTORE fails to recover NIAH while degrading MATH500, confirming that $\mathbf{W}_V$ does not contribute to routing recovery. Q-RESTORE and K-RESTORE each achieve partial recall recovery but fall significantly short at 128K (e.g., K-RESTORE collapses on NIAH-Single-3 (13.2 vs. 30.8), indicating that isolated restoration introduces a $\mathbf{W}_Q$ - $\mathbf{W}_K$ mismatch and that joint restoration is necessary for coherent routing geometry. QK-PRO achieves the highest performance on MATH500, but trading long-range routing recovery for math preservation. QK-FROZEN consistently underperforms QK-RESTORE in both NIAH and mathematical tasks, indicating that unconstrained SFT is for better $\mathbf{W}_V$ adaptation.

8 Conclusion

We identify CoT-SFT-induced long-context recall degradation in distilled hybrid models. Reasoning-focused SFT biases attention toward short-range interactions, distorting query-key routing in softmax-attention layers. We propose QK-RESTORE, a training-free method that restores only query and key projections from the pre-SFT checkpoint. Across architectures, QK-RESTORE substantially recovers long-context recall while preserving reasoning gains, highlighting routing stability as key to efficient long-context models.

Limitation

First, all experiments are conducted after CoT math fine-tuning; it remains open whether the intervention generalises to the general instruction-tuning setting, where training mixes short-answer, multi-turn dialogue, and retrieval-heavy formats. Second, QK-RESTORE currently treats all softmax-attention layers and all heads uniformly — it restores everything or nothing. In practice, not every head loses its long-range routing equally: some heads may collapse completely while others remain intact. We leave them as the exciting directions for future work.

Acknowledgment

Zhijiang Guo is supported by the Guangdong Provincial Key Lab of Integrated Communication, Sensing and Computation for Ubiquitous Internet of Things (No. 2023B1212010007).

References

- Loubna Ben Allal, Anton Lozhkov, Elie Bakouch, Gabriel Martín Blázquez, Guilherme Penedo, Lewis Tunstall, Andrés Marafioti, Hynek Kydlíček, Agustín Piqueres Lajarín, Vaibhav Srivastav, Joshua Lochner, Caleb Fahlgren, Xuan-Son Nguyen, Clémentine Fourrier, Ben Burtenshaw, Hugo Larcher, Haojun Zhao, Cyril Zakka, Mathieu Morlon, and 3 others. 2025. [Smollm2: When smol goes big – data-centric training of a small language model](#). *Preprint*, arXiv:2502.02737.
- Simran Arora, Brandon Yang, Sabri Eyuboglu, Avanika Narayan, Andrew Hojel, Immanuel Trummer, and Christopher Ré. 2023. Language models enable simple systems for generating structured views of heterogeneous data lakes. *arXiv preprint arXiv:2304.09433*.
- Yingfa Chen, Zhen Leng Thai, Zihan Zhou, Zhu Zhang, Xingyu Shen, Shuo Wang, Chaojun Xiao, Xu Han, and Zhiyuan Liu. 2026. [Hybrid linear attention done right: Efficient distillation and effective architectures for extremely long contexts](#). *Preprint*, arXiv:2601.22156.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. 2021. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*.
- Tri Dao and Albert Gu. 2024. [Transformers are ssms: Generalized models and efficient algorithms through structured state space duality](#). *Preprint*, arXiv:2405.21060.
- Leo Gao, Jonathan Tow, Baber Abbasi, Stella Biderman, Sid Black, Anthony DiPofi, Charles Foster, Laurence Golding, Jeffrey Hsu, Alain Le Noac’h, Haonan Li, Kyle McDonell, Niklas Muennighoff, Chris Ociepa, Jason Phang, Laria Reynolds, Hailey Schoelkopf, Aviya Skowron, Lintang Sutawika, and 5 others. 2024. [The language model evaluation harness](#).
- Daniel Goldstein, Eric Alcaide, Janna Lu, and Eugene Cheah. 2026. [Radlads: Rapid attention distillation to linear attention decoders at scale](#). *Preprint*, arXiv:2505.03005.
- Albert Gu and Tri Dao. 2024. [Mamba: Linear-time sequence modeling with selective state spaces](#). *Preprint*, arXiv:2312.00752.
- Yuxian Gu, Qinghao Hu, Haocheng Xi, Junyu Chen, Shang Yang, Song Han, and Han Cai. 2026. Jet-nemotron: Efficient language model with post neural architecture search. *Advances in Neural Information Processing Systems*, 38:47191–47218.
- Etash Guha, Ryan Marten, Sedrick Keh, Negin Raoof, Georgios Smyrnis, Hritik Bansal, Marianna Nezhurina, Jean Mercat, Trung Vu, Zayne Sprague, and 1 others. 2025. Openthoughts: Data recipes for reasoning models. *arXiv preprint arXiv:2506.04178*.
- Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Peiyi Wang, Qihao Zhu, Runxin Xu, Ruoyu Zhang, Shirong Ma, Xiao Bi, Xiaokang Zhang, Xingkai Yu, Yu Wu, Z. F. Wu, Zhibin Gou, Zhihong Shao, Zhuoshu Li, Ziyi Gao, Aixin Liu, and 175 others. 2025. [Deepseek-r1 incentivizes reasoning in llms through reinforcement learning](#). *Nature*, 645(8081):633–638.
- Yuichiro Hoshino, Hideyuki Tachibana, Muneyoshi Inahara, and Hiroto Takegawa. 2025. [Rad: Redundancy-aware distillation for hybrid models via self-speculative decoding](#). *Preprint*, arXiv:2505.22135.
- Cheng-Ping Hsieh, Simeng Sun, Samuel Krizan, Shantanu Acharya, Dima Rekes, Fei Jia, Yang Zhang, and Boris Ginsburg. 2024. Ruler: What’s the real context size of your long-context language models? *arXiv preprint arXiv:2404.06654*.
- Dong Huang, Guangtao Zeng, Jianbo Dai, Meng Luo, Han Weng, Yuhao Qing, Heming Cui, Zhijiang Guo, and Jie M Zhang. 2024. Efficoder: Enhancing code generation in large language models through efficiency-aware fine-tuning. *arXiv preprint arXiv:2410.10209*.
- Hugging Face. 2025. [Open r1: A fully open reproduction of deepseek-r1](#).
- Samy Jelassi, David Brandfonbrener, Sham M Kakade, and Eran Malach. 2024. Repeat after me: Transformers are better than state space models at copying. *arXiv preprint arXiv:2402.01032*.

- Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, L  lio Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timoth  e Lacroix, and William El Sayed. 2023. [Mistral 7b](#). *Preprint*, arXiv:2310.06825.
- Angelos Katharopoulos, Apoorv Vyas, Nikolaos Pappas, and Fran  ois Fleuret. 2020. [Transformers are rnn: Fast autoregressive transformers with linear attention](#). *Preprint*, arXiv:2006.16236.
- Xingxuan Li, Yao Xiao, Dianwen Ng, Hai Ye, Yue Deng, Xiang Lin, Bin Wang, Zhanfeng Mo, Chong Zhang, Yueyi Zhang, Zonglin Yang, Ruilin Li, Lei Lei, Shihao Xu, Han Zhao, Weiling Chen, Feng Ji, and Lidong Bing. 2025a. [Miromind-m1: An open-source advancement in mathematical reasoning via context-aware multi-stage policy optimization](#). *Preprint*, arXiv:2507.14683.
- Yanhong Li, Songlin Yang, Shawn Tan, Mayank Mishra, Rameswar Panda, Jiawei Zhou, and Yoon Kim. 2025b. Distilling to hybrid attention models via kl-guided layer selection. *arXiv preprint arXiv:2512.20569*.
- Hunter Lightman, Vineet Kosaraju, Yura Burda, Harri Edwards, Bowen Baker, Teddy Lee, Jan Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. 2023. [Let’s verify step by step](#). *Preprint*, arXiv:2305.20050.
- Colin Lockard, Prashant Shiralkar, and Xin Luna Dong. 2019. Openceres: When open information extraction meets the semi-structured web. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 3047–3056.
- Anton Lozhkov, Loubna Ben Allal, Leandro von Werra, and Thomas Wolf. 2024. [Fineweb-edu: the finest collection of educational content](#).
- Justus Matternt, Sami Jaghouar, Manveer Basra, Jan-nik Straube, Matthew Di Ferrante, Felix Gabriel, Jack Min Ong, Vincent Weisser, and Johannes Hagemann. 2025. [Synthetic-1: Two million collaboratively generated reasoning traces from deepseek-r1](#).
- Arindam Mitra, Hamed Khanpour, Corby Rosset, and Ahmed Awadallah. 2024. [Orca-math: Unlocking the potential of slms in grade school math](#). *Preprint*, arXiv:2402.14830.
- NVIDIA, :, Aaron Blakeman, Aaron Grattafiori, Aarti Basant, Abhibha Gupta, Abhinav Khattar, Adi Renduchintala, Aditya Vavre, Akanksha Shukla, Akhiad Bercovich, Aleksander Ficek, Aleksandr Shaposhnikov, Alex Kondratenko, Alexander Bukharin, Alexandre Milesi, Ali Taghibakhshi, Alisa Liu, Amelia Barton, and 340 others. 2025. [Nvidia nemotron 3: Efficient and open intelligence](#). *Preprint*, arXiv:2512.20856.
- Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul Christiano, Jan Leike, and Ryan Lowe. 2022. [Training language models to follow instructions with human feedback](#). *Preprint*, arXiv:2203.02155.
- Ben Prystawski, Michael Y. Li, and Noah D. Goodman. 2023. [Why think step by step? reasoning emerges from the locality of experience](#). *Preprint*, arXiv:2304.03843.
- Zhen Qin, Weigao Sun, Dong Li, Xuyang Shen, Weixuan Sun, and Yiran Zhong. 2024a. [Lightning attention-2: A free lunch for handling unlimited sequence lengths in large language models](#). *Preprint*, arXiv:2401.04658.
- Zhen Qin, Weigao Sun, Dong Li, Xuyang Shen, Weixuan Sun, and Yiran Zhong. 2024b. [Various lengths, constant speed: Efficient language modeling with lightning attention](#). *Preprint*, arXiv:2405.17381.
- Qwen, :, An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, Huan Lin, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiaxi Yang, Jingren Zhou, and 25 others. 2025. [Qwen2.5 technical report](#). *Preprint*, arXiv:2412.15115.
- Qwen Team. 2026. [Qwen3.5: Towards native multi-modal agents](#).
- Kimi Team, Yu Zhang, Zongyu Lin, Xingcheng Yao, Jiaxi Hu, Fanqing Meng, Chengyin Liu, Xin Men, Songlin Yang, Zhiyuan Li, Wentao Li, Enzhe Lu, Weizhou Liu, Yanru Chen, Weixin Xu, Longhui Yu, Yejie Wang, Yu Fan, Longguang Zhong, and 41 others. 2025. [Kimi linear: An expressive, efficient attention architecture](#). *Preprint*, arXiv:2510.26692.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2023. [Attention is all you need](#). *Preprint*, arXiv:1706.03762.
- Dustin Wang, Rui-Jie Zhu, Steven Abreu, Yong Shan, Taylor Kergan, Yuqi Pan, Yuhong Chou, Zheng Li, Ge Zhang, Wenhao Huang, and Jason Eshraghian. 2025. [A systematic analysis of hybrid linear attention](#). *Preprint*, arXiv:2507.06457.
- Zihan Wang, Yijun Dong, and Qi Lei. 2026. [When does chain-of-thought help: A markovian perspective](#). *Preprint*, arXiv:2603.00306.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed Chi, Quoc Le, and Denny Zhou. 2023. [Chain-of-thought prompting elicits reasoning in large language models](#). *Preprint*, arXiv:2201.11903.

Liang Wen, Yunke Cai, Fenrui Xiao, Xin He, Qi An, Zhenyu Duan, Yimin Du, Junchen Liu, Lifu Tang, Xiaowei Lv, Haosheng Zou, Yongchao Deng, Shousheng Jia, and Xiangzheng Zhang. 2025. [Light-rl: Curriculum sft, dpo and rl for long cot from scratch and beyond](#). *Preprint*, arXiv:2503.10460.

Haotian Xu, Xing Wu, Weinong Wang, Zhongzhi Li, Da Zheng, Boyuan Chen, Yi Hu, Shijia Kang, Jiaming Ji, Yingying Zhang, and 1 others. 2025. Redstar: does scaling long-cot data unlock better slow-reasoning systems? *arXiv preprint arXiv:2501.11284*.

Mingyu Yang, Mehdi Rezagholizadeh, Guihong Li, Vikram Appia, and Emad Barsoum. 2025. Zebra-llama: Towards extremely efficient hybrid models. *arXiv preprint arXiv:2505.17272*.

Songlin Yang, Bailin Wang, Yikang Shen, Rameswar Panda, and Yoon Kim. 2024. [Gated linear attention transformers with hardware-efficient training](#). *Preprint*, arXiv:2312.06635.

Qiyang Yu, Zheng Zhang, Ruofei Zhu, Yufeng Yuan, Xiaochen Zuo, Yu Yue, Weinan Dai, Tiantian Fan, Gaohong Liu, Lingjun Liu, Xin Liu, Haibin Lin, Zhiqi Lin, Bole Ma, Guangming Sheng, Yuxuan Tong, Chi Zhang, Mofan Zhang, Wang Zhang, and 16 others. 2025. [Dapo: An open-source llm reinforcement learning system at scale](#). *Preprint*, arXiv:2503.14476.

Michael Zhang, Kush Bhatia, Hermann Kumbong, and Christopher Ré. 2024. [The hedgehog & the porcupine: Expressive linear attentions with softmax mimicry](#). *Preprint*, arXiv:2402.04347.

Yifan Zhang and Team Math-AI. 2024. American invitational mathematics examination (aime) 2024.

Chuyue Zhou, Hongya Lyu, Xinle Lin, Hengyu Zhao, Junshao Guo, Xueren Zhang, Shuaikang Xue, Qiang Ma, Jie Zhou, Yudong Wang, and Zhiyuan Liu. 2026. [Ultradata-math](#).

A Extended Related Work

Efficient Long-Context Models and Hybrid Attention. Recent hybrid attention models interleave softmax-attention layers with efficient recurrent or linear-attention layers, achieving competitive performance on both commonsense reasoning and recall-intensive long-context tasks while improving inference efficiency (Qwen Team, 2026; Team et al., 2025; NVIDIA et al., 2025). Pretraining strong hybrid models from scratch at large scale remains prohibitively expensive, which motivates recent efforts to obtain hybrid models through Transformer-to-hybrid conversion or distillation (Chen et al., 2026; Li et al., 2025b). However,

existing work primarily focus to obtaining high-performing converted hybrid checkpoints, while their behavior during subsequent post-training remains underexplored.

Distilling Transformers into Hybrid Models.

Transformer-to-hybrid distillation converts selected softmax-attention layers into linear or recurrent mixers, where layer selection critically influences long-context retrieval performance (Goldstein et al., 2026; Chen et al., 2026; Li et al., 2025b; Gu et al., 2026). However, strong recall performance after conversion does not necessarily imply stability after reasoning-oriented post-training; specifically, across multiple distilled hybrid models with strong long-context retrieval, we observe that CoT-SFT can substantially degrade recall capabilities, revealing a critical failure mode that remains unaddressed by existing distillation methods.

B Discussion of the CoT-Markov Structure Assumption

Discussion 1. The 1-step Markov condition is an idealisation for analysis. Any k -step Markov chain is equivalent to a 1-step chain on a K^k -state product space, yielding the same exponential decay with a larger reasoning horizon W .

Discussion 2. Reversibility is invoked in the proof of Lemma 4.5 at one step: it implies that the Markov operator P is self-adjoint on $L^2(\pi)$. We acknowledge that a causally generated sequence such as a CoT reasoning chain does not, in general, satisfy detailed balance: the latent transition probabilities are asymmetric by construction. Reversibility is, however, a *sufficient* condition rather than a necessary one. For a general ergodic chain (possibly non-reversible), the identical bound holds with ρ replaced by the second-largest *singular value* of P , obtained via its singular value decomposition in place of the eigendecomposition. Since singular values are always real and non-negative, the proof proceeds without modification, and the form of the final bound is unchanged. We therefore state Lemma 4.5 under reversibility for clarity of exposition, with the understanding that all subsequent results hold for the general ergodic case by replacing eigenvalues with singular values throughout.

C Justification for Norm Bounds Assumption

Bound on $\|v_s\|$. The value vector is $v_s = W_V \text{LN}(h_s)$, where LN denotes layer normalisation.

tion. Layer normalisation standardises each coordinate so that $\text{LN}(\mathbf{h}_s)_i = \gamma_i(\mathbf{h}_{s,i} - \hat{\mu})/\hat{\sigma} + \beta_i$, where $\hat{\mu}$ and $\hat{\sigma}$ are the empirical mean and standard deviation of the entries of $\mathbf{h}_s$. After normalisation, the pre-affine vector satisfies $\sum_i z_i^2 = d$, giving

$$\|\text{LN}(\mathbf{h}_s)\|^2 = \sum_i (\gamma_i z_i + \beta_i)^2 \quad (19)$$

$$\leq 2(\|\gamma\|_\infty^2 d + \|\beta\|^2) \quad (20)$$

$$\leq C_{\text{LN}}^2 d, \quad (21)$$

for a constant C_{LN} depending only on the learned scale γ and bias β . The projection matrix $\mathbf{W}_V$ is a fixed trained weight; its operator norm $\|\mathbf{W}_V\|_{\text{op}} < \infty$ is a finite model-dependent constant. Combining we have:

$$\|\mathbf{v}_s\| = \|\mathbf{W}_V \text{LN}(\mathbf{h}_s)\| \quad (22)$$

$$\leq \|\mathbf{W}_V\|_{\text{op}} C_{\text{LN}} \sqrt{d} =: B_V. \quad (23)$$

Bound on $\mathbf{G}_t$. The gradient $\mathbf{G}_t = \partial \mathcal{L} / \partial \mathbf{o}_t$ back-propagates through the final layer normalisation (with Jacobian $\mathbf{J}_{\text{LN}}$) and the language-model head $\mathbf{W}_{\text{lm}}$, giving

$$\mathbf{G}_t = \mathbf{J}_{\text{LN}}^T \mathbf{W}_{\text{lm}}^T (\mathbf{p}_t - \mathbf{y}_t), \quad (24)$$

where $\mathbf{p}_t \in \Delta^{|\mathcal{V}|-1}$ is the predicted distribution and $\mathbf{y}_t \in \{0, 1\}^{|\mathcal{V}|}$ is the one-hot target. By submultiplicativity of the operator norm:

$$\|\mathbf{G}_t\| \leq \|\mathbf{J}_{\text{LN}}\|_{\text{op}} \|\mathbf{W}_{\text{lm}}\|_{\text{op}} \|\mathbf{p}_t - \mathbf{y}_t\|. \quad (25)$$

Since $\mathbf{p}_t - \mathbf{y}_t$ is a difference of probability vectors, $\|\mathbf{p}_t - \mathbf{y}_t\| \leq \sqrt{2}$. For a fixed trained model, both $\|\mathbf{J}_{\text{LN}}\|_{\text{op}}$ and $\|\mathbf{W}_{\text{lm}}\|_{\text{op}}$ are finite constants. Therefore:

$$\|\mathbf{G}_t\| \leq \sqrt{2} \|\mathbf{J}_{\text{LN}}\|_{\text{op}} \|\mathbf{W}_{\text{lm}}\|_{\text{op}} =: B_G. \quad (26)$$

D Justification for Score-function Identity Assumption

At the optimum $\mathbf{p}_\theta = \mathbf{p}_{\text{data}}$, the cross-entropy loss with respect to the attention output $\mathbf{o}_t$ satisfies, for any fixed context $x_{1:t}$:

$$\mathbb{E}_{y_t} \left[\frac{\partial}{\partial \mathbf{o}_t} (-\log p_\theta(y_t | x_{1:t})) \right] \quad (27)$$

$$= -\frac{\partial}{\partial \mathbf{o}_t} \sum_{y_t} p(y_t | x_{1:t}) \log p_\theta(y_t | x_{1:t}) \Big|_{\mathbf{p}_\theta = \mathbf{p}_{\text{data}}} \quad (28)$$

$$= 0, \quad (29)$$

since the cross-entropy gradient vanishes at its own minimum. Applying the law of total expectation over $x_{1:t}$ then yields $\mathbb{E}[\mathbf{G}_t] = 0$. More precisely, because the gradient is computed at a given context $x_{1:t}$ and then averaged over the target y_t drawn from the conditional, the identity holds conditionally:

$$\mathbb{E}[\mathbf{G}_t | x_{1:t}] = 0. \quad (30)$$

E Preliminary Definition and Lemma

We list the following necessary definitions and lemmas for the proof.

Definition E.1 ($L^2(\pi)$ Space). For a Markov chain with finite state space $\mathcal{Z}$ and stationary distribution π , define the weighted function space:

$$L^2(\pi) = \{f : \mathcal{Z} \rightarrow \mathbb{R}\} \quad (31)$$

with inner product

$$\langle f, g \rangle_\pi := \sum_{z \in \mathcal{Z}} \pi(z) f(z) g(z) \quad (32)$$

$$= \mathbb{E}_\pi[f(Z) g(Z)] \quad (33)$$

and norm $\|f\|_{L^2(\pi)} := \sqrt{\langle f, f \rangle_\pi}$.

In addition, mean-zero means $\langle f, 1 \rangle_\pi = \mathbb{E}_\pi[f(Z)] = 0$, i.e., the function has zero expectation under π .

Definition E.2 (Operator P acting on functions). We define the Markov operator P acting on functions:

$$(Pf)(z) := \sum_{z'} P(z, z') f(z') \quad (34)$$

$$= \mathbb{E}[f(Z_{t+1}) | Z_t = z] \quad (35)$$

Lemma E.3. For a reversible chain, detailed balance holds: $\pi(z)P(z, z') = \pi(z')P(z', z)$, which makes P symmetric in $L^2(\pi)$.

Proof. According to the definition of Markov operator P in E.2, we have:

$$\langle Pf, g \rangle_\pi = \sum_z \pi(z) (Pf)(z) g(z) \quad (36)$$

$$= \sum_{z, z'} \pi(z) P(z, z') f(z') g(z) \quad (37)$$

$$= \sum_{z, z'} \pi(z') P(z', z) f(z') g(z) \quad (38)$$

$$= \sum_{z'} \pi(z') f(z') \sum_z P(z', z) g(z) \quad (39)$$

$$= \langle f, Pg \rangle_\pi \quad (40)$$

□

Since $|\mathcal{Z}| = K < \infty$, the finite-dimensional spectral theorem for self-adjoint operators yields a real orthonormal eigenbasis $\{\varphi_k\}_{k=1}^K$ of $L^2(\pi)$ with eigenvalues $1 = \lambda_1 \geq |\lambda_2| \geq \dots \geq |\lambda_K|$, and spectral gap $1 - \rho$ where $\rho := \max_{k \geq 2} |\lambda_k|$:

$$P\varphi_k = \lambda_k \varphi_k, \langle \varphi_k, \varphi_j \rangle_\pi = \delta_{kj} \quad (41)$$

Further, for mean-zero f :

$$\langle f, \varphi_1 \rangle_\pi = \langle f, \mathbf{1} \rangle_\pi = \mathbb{E}_\pi[f(Z)] = 0 \quad (42)$$

which means f has no component along ψ_1 . Its expansion uses only $k \geq 2$:

$$f = \sum_{k \geq 2} a_k \varphi_k, \quad a_k = \langle f, \varphi_k \rangle_\pi \quad (43)$$

F Proof for Spectral Correlation Decay Lemma

Proof. At stationarity ($z_s \sim \pi$), the Tower Property tells:

$$\mathbb{E}[f(z_t)g(z_s)] = \mathbb{E}[\mathbb{E}[f(z_t) | z_s] \cdot g(z_s)] \quad (44)$$

Note that the inner expectation is the τ -step prediction: for a fixed starting state $z_s = z$, the expected value of $f(z_t)$ after $\tau = t - s$ steps is exactly $(P^\tau f)(z)$:

$$(P^\tau f)(z) = \mathbb{E}[f(z_t) | z_s = z] \quad (45)$$

Therefore:

$$\mathbb{E}[f(z_t)g(z_s)] = \mathbb{E}[(P^\tau f)(z_s)g(z_s)] \quad (46)$$

$$= \sum_z \pi(z) (P^\tau f)(z) g(z) \quad (47)$$

$$= \langle P^\tau f, g \rangle_\pi \quad (48)$$

Then we expand f, g by the orthonormal eigenbasis:

$$f = \sum_{k \geq 2} a_k \varphi_k, g = \sum_{k \geq 2} b_k \varphi_k, \quad (49)$$

Using the orthonormality $\langle \varphi_k, \varphi_j \rangle_\pi = \delta_{kj}$, we further simplify $\langle P^\tau f, g \rangle_\pi$:

$$\langle P^\tau f, g \rangle_\pi = \left\langle \sum_{k \geq 2} a_k \lambda_k^\tau \varphi_k, \sum_{j \geq 2} b_j \varphi_j \right\rangle_\pi \quad (50)$$

$$= \sum_{k \geq 2} a_k \lambda_k^\tau b_k \quad (51)$$

Then we bound it:

$$|\langle P^\tau f, g \rangle_\pi| = \left| \sum_{k \geq 2} a_k \lambda_k^\tau b_k \right| \quad (52)$$

$$\leq \sum_{k \geq 2} |a_k| |\lambda_k|^\tau |b_k| \quad (53)$$

$$\leq \rho^\tau \sum_{k \geq 2} |a_k| |b_k| \quad (54)$$

Applying the Cauchy-Schwarz Inequality:

$$\sum_{k \geq 2} |a_k| |b_k| \leq \sqrt{\sum_{k \geq 2} |a_k|^2} \sqrt{\sum_{k \geq 2} |b_k|^2} \quad (55)$$

$$= \|f\|_{L^2(\pi)} \|g\|_{L^2(\pi)} \quad (56)$$

Now combining all together:

$$|\mathbb{E}[f(z_t)g(z_s)]| \leq \rho^\tau \|f\|_{L^2(\pi)} \|g\|_{L^2(\pi)} \quad (57)$$

$$= \rho^{t-s} \|f\|_{L^2(\pi)} \|g\|_{L^2(\pi)} \quad (58)$$

□

G Proof of the Gradient Locality Theorem

Proof. For clarity, we drop the layer/head superscripts in the following derivation. First, we bound the gradient:

$$\left| \frac{\partial \mathcal{L}}{\partial e_{ts}} \right| \leq |\mathbf{G}_t \cdot \mathbf{v}_s| + |\mathbf{G}_t \cdot \bar{\mathbf{v}}_t| \quad (59)$$

Then we consider how to bound $\mathbb{E}[|\mathbf{G}_t \cdot \mathbf{v}_s|]$. Applying the Tower Property:

$$\mathbb{E}[\mathbf{G}_t \cdot \mathbf{v}_s] = \mathbb{E}[\mathbb{E}[\mathbf{G}_t | z_t] \cdot \mathbb{E}[\mathbf{v}_s | z_s]] \quad (60)$$

$$= \mathbb{E}[\phi(z_t) \cdot \psi(z_s)] \quad (61)$$

where we define:

$$\phi(z) = \mathbb{E}[\mathbf{G}_t | z_t = z], \psi(z) = \mathbb{E}[\mathbf{v}_s | z_s = z] \quad (62)$$

According to Assumption 4.4:

$$\mathbb{E}_\pi[\phi(z)] = \sum_z \pi(z) \phi(z) \quad (63)$$

$$= \sum_z P(z_t = z) \phi(z) \quad (64)$$

$$= \mathbb{E}[\phi(z_t)] = \mathbb{E}[\mathbb{E}[\mathbf{G}_t | z_t = z]] \quad (65)$$

$$= \mathbb{E}[\mathbf{G}_t] = \mathbb{E}[\mathbf{G}_t | x_{1:t}] = \mathbf{0} \quad (66)$$

Similarly, centering ψ does not change the covariance (since $\mathbb{E}[\phi(z_t)] = 0$). Both ϕ and ψ are mean-zero functions on $\mathcal{Z}$.

Then applying Lemma 4.5:

$$|\mathbb{E}[\mathbf{G}_t \cdot \mathbf{v}_s]| = |\mathbb{E}[\phi(z_t) \cdot \psi(z_s)]| \quad (67)$$

$$\leq \rho^\tau \|\phi\|_{L^2(\pi)} \|\psi\|_{L^2(\pi)} \quad (68)$$

$$\leq B_G B_V \rho^\tau \quad (69)$$

Applying the same bound to the $\bar{\mathbf{v}}_t$ term, which satisfies $\|-\bar{\mathbf{v}}_t\| \leq B_V$ as a convex combination:

$$g(\tau) = \mathbb{E} \left[\left| \frac{\partial \mathcal{L}}{\partial e_{ts}} \right| \right] \leq 2B_G B_V \rho^\tau = C' e^{-\tau/W}, \quad (70)$$

where $C' = 2B_G B_V$ and $W = \frac{-1}{\log \rho}$. $\square$

H Proof for Token Autocorrelation Decay

Define the centred indicator feature $f_v(x) := \mathbf{1}[x = v] - \pi(v)$ and the π -weighted autocorrelation at lag τ :

$$\bar{\rho}(\tau) := \sum_v \pi(v) \frac{\text{Cov}(f_v(x_t), f_v(x_{t+\tau}))}{\text{Var}(f_v(x_t))} \quad (71)$$

For the numerator $\text{Cov}(f_v(x_t), f_v(x_{t+\tau}))$, we expand it as:

$$\mathbb{E}[f_v(x_t) f_v(x_{t+\tau})] - \mathbb{E}[f_v(x_t)] \mathbb{E}[f_v(x_{t+\tau})] \quad (72)$$

Considering the expectation of $f_v(x)$:

$$\mathbb{E}[f_v(x_t)] = \mathbb{E}[\mathbf{1}[x_t = v] - \pi(v)] \quad (73)$$

$$= \pi(v) - \pi(v) = 0 \quad (74)$$

Therefore we can further simplify:

$$\bar{\rho}(\tau) = \sum_v \pi(v) \frac{\mathbb{E}[f_v(x_t) f_v(x_{t+\tau})]}{\pi(v)(1 - \pi(v))} \quad (75)$$

In the HMM, tokens are conditionally independent given their latent states: $x_t \perp x_{t+\tau} \mid z_t, z_{t+\tau}$. By the tower property:

$$\mathbb{E}[f_v(x_t) f_v(x_{t+\tau})] \quad (76)$$

$$= \mathbb{E}[\mathbb{E}[f_v(x) \mid z_t] \cdot \mathbb{E}[f_v(x) \mid z_{t+\tau}]] \quad (77)$$

$$= \mathbb{E}[\phi_v(z_t) \cdot \phi_v(z_{t+\tau})] \quad (78)$$

where $\phi_v(z) = \mathbb{E}[f_v(x)|z] = P(x = v|z) - \pi(v)$ is a mean-zero function of z :

$$\mathbb{E}_\pi[\phi_v(z)] = \mathbb{E}_\pi[P(x = v|z) - \pi(v)] \quad (79)$$

$$= \sum_z \pi(z) P(x = v|z) - \pi(v) \quad (80)$$

$$= \pi(v) - \pi(v) = 0 \quad (81)$$

Since both $\phi_v(z_t)$ and $\phi_v(z_{t+\tau})$ are mean-zero, we can apply Lemma 4.5 directly:

$$|\mathbb{E}[\phi_v(z_t) \cdot \phi_v(z_{t+\tau})]| \leq \rho^\tau \|\phi_v\|_{L^2(\pi)}^2 \quad (82)$$

where $\rho = \lambda_2$ is the second-largest eigenvalue of the latent transition matrix P . Now we can bound the autocorrelation:

$$\bar{\rho}(\tau) = \sum_v \pi(v) \frac{\mathbb{E}[f_v(x_t) f_v(x_{t+\tau})]}{\pi(v)(1 - \pi(v))} \quad (83)$$

$$\leq \rho^\tau \sum_v \frac{\|\phi_v\|_{L^2(\pi)}^2}{\pi(v)(1 - \pi(v))} = C \cdot \rho^\tau \quad (84)$$

Conclusion. Observing $\bar{\rho}(\tau) \sim C \hat{\rho}^\tau$ with $\hat{\rho} < 1$ in the corpus confirms that the latent chain has spectral gap $1 - \hat{\rho} > 0$, which is exactly Lemma 4.5's hypothesis.

I Discussion of Attention Gradient Decay in Pre-training Stage

Theorem 4.6 rests on Lemma 4.5, which holds for any ergodic Markov chain with a spectral gap: a property of the data-generating process, not specific to CoT structure, which is also observed in general pre-training data corpus.

Consequently, the gradient locality bound $g(\tau) \leq C' e^{-\tau/W}$ applies to any training corpus, including general pre-training data. The severity of routing collapse is governed by the mismatch $\Delta W = W_{\text{corr}} - W_{\text{grad}}$: how far the data demands the model to attend vs. how far the gradient actually reinforces it. For general pre-training corpora, this mismatch is smaller than for CoT, so routing degrades more slowly but by the same mechanism.

For hybrid architectures, however, even the slower pre-training erosion meaningfully reduces available recall capacity before fine-tuning begins, leaving the model closer to the threshold. The hybrid architecture's vulnerability therefore originates during pre-training and is further amplified by CoT fine-tuning.

J Proof for the Routing-Extraction Gradient Decoupling Theorem

Proof. From Equation 7, we have:

$$\nabla_{\mathbf{W}^{Q,(h)}} \mathcal{L} = \frac{1}{\sqrt{d_h}} \sum_{t,s} \frac{\partial \mathcal{L}}{\partial e_{ts}^{(h)}} \mathbf{k}_s^{(h')} \mathbf{h}_t^T \quad (85)$$

where h' is the corresponding KV head. Then considering its Frobenius norm:

$$\mathbb{E} \left[\left\| \frac{\partial \mathcal{L}}{\partial e_{ts}} \mathbf{k}_s \mathbf{h}_t^T \right\|_F \right] = \mathbb{E} \left[\left\| \frac{\partial \mathcal{L}}{\partial e_{ts}} \right\| \right] \cdot \|\mathbf{k}_s\| \cdot \|\mathbf{h}_t\| \quad (86)$$

$$\leq C' B_K B_h \rho^\tau \quad (87)$$

$$= C_R \cdot \rho^\tau \quad (88)$$

via Theorem 4.6 and bounded norms $\|\mathbf{k}_s\| \leq B_K$, $\|\mathbf{h}_t\| \leq B_h$.

For the Extraction Parameters, the value vector $\mathbf{v}_s = \mathbf{W}_V \mathbf{h}_s$ enters the loss through $\mathbf{o}_t = \sum_s A_{ts} \mathbf{v}_s$. Accumulating all downstream dependencies:

$$\frac{\partial \mathcal{L}}{\partial \mathbf{v}_s} = \sum_{t \geq s} A_{ts} \mathbf{G}_t \quad (89)$$

Now we assume that $A_{ss} \geq \delta_A > 0$ for all s and $\mathbb{E}[\|\mathbf{G}_s\|] \geq c_G > 0$ for all s :

$$\mathbb{E} \left[\left\| \frac{\partial \mathcal{L}}{\partial \mathbf{v}_s} \right\| \right] \geq \mathbb{E}[A_{ss} \|\mathbf{G}_s\|] \geq \delta_A c_G > 0 \quad (90)$$

□

K Training Data and Configuration Details

For HypeNet, we conduct both the pre-training and SFT.

Pre-training Data. Since we mainly focus on the long-context recall and math reasoning ability, during the pre-training stage, we mix the FineWeb-Edu (Lozhkov et al., 2024) and UltraData-Math (Zhou et al., 2026). Following observations from recent pretraining studies that maintain web corpora as the dominant source while increasing structured reasoning data (Allal et al., 2025), we construct a mixture with 80% general web data and 20% math-focused data. Prior evidence indicates that web-scale corpora preserve broad linguistic and factual competence, while specialized mathematical corpora improve reasoning and STEM performance.

Pre-training Config. We adopt the 3-stage training for HypeNet, keeping the same setting in Chen et al. (2026), shown in Table 5, Table 6, and Table 7.

SFT Data. To further improve the model’s math reasoning performance, we conduct CoT-SFT with MiroMind-M1 dataset (Li et al., 2025a), which is collected from OpenR1 (Hugging Face, 2025), Open-thoughts (Guha et al., 2025), Light-R1 (Wen et al., 2025), and Synthetic-1 (Mattern et al., 2025).

SFT Config. The detailed configurations are shown in Table 8 and Table 9.

L Evaluation Details

For evaluating NIAH, we apply LM Evaluation Harness² (Gao et al., 2024) for official test. For evaluating MATH500 and GSM8K, we report average pass@1 over 16 independent generations and Maj@16 (i.e., majority vote as the final prediction) per problem. In detail, for MATH500, the generation length is 8,192; for GSM8K, we set it as 2,048, since it is easier. For LiveCodeBench, we report the average pass@1 over 8 independent generations with 16,384 response length.

M Analysis on Pure Softmax-Attention Model

Our main experiments concentrate on the hybrid models, and in this section, we explore whether SFT can degrade the pure softmax-attention models as well. We include Qwen2.5-3B, Qwen2.5-7B (Qwen et al., 2025) and Mistral-7B-Instruct-v0.3 (Jiang et al., 2023) for investigation. We include the details on how we identify the top layers for retrieval in Appendix N.

Analysis. Table 10 shows that for Mistral-7B, it exhibits strong sensitivity to SFT at longer contexts. SFT causes severe degradation on NIAH-Single-3 at 64K, dropping from 81.0 to 4.80, suggesting a breakdown of long-range retrieval. QK-RESTORE recovers this degradation limitedly, indicating weaker long-context representations in the pre-trained model.




Figure 7: Attention gradient decay on Qwen2.5-3B. We observe a different pattern from the exponential decay of HypeNet.

For Qwen2.5-7B and Qwen2.5-3B, long-context performance is highly stable across NIAH-Single-2 and Single-3. SFT induces only minor fluctuations

²<https://github.com/EleutherAI/lm-evaluation-harness>

Table 5: Pre-Training Conigs for HypeNet-2B

STAGE	TOKENS	LR	LR SCHEDULER	CONTEXT LEN.	BATCH	TRAINING STEPS
1	320M	1e-3 $\rightarrow$ 1e-5	Cosine	512	32	20,000
2	1B	1e-4 $\rightarrow$ 1e-5	Cosine	512	96	20,000
3	1B	1e-5	Constant	16,384	128	500

Table 6: Pre-Training Conigs for HypeNet-5B

STAGE	TOKENS	LR	LR SCHEDULER	CONTEXT LEN.	BATCH	TRAINING STEPS
1	320M	1e-3 $\rightarrow$ 1e-5	Cosine	512	32	20,000
2	1B	5e-5 $\rightarrow$ 1e-5	Cosine	512	96	20,000
3	1B	1e-5	Constant	16,384	128	500

Table 7: Pre-Training Conigs for HypeNet-9B

STAGE	TOKENS	LR	LR SCHEDULER	CONTEXT LEN.	BATCH	TRAINING STEPS
1	320M	1e-3 $\rightarrow$ 1e-5	Cosine	512	32	20,000
2	1B	3e-5 $\rightarrow$ 1e-5	Cosine	512	96	20,000
3	1B	1e-5	Constant	16,384	128	500

Table 8: SFT Conigs for 2B scale model

LR	LR SCHEDULER	CONTEXT LEN.	BATCH	TRAINING STEPS
1e-5	Constant	16,384	128	100

Table 9: SFT Conigs for 5B, 9B scale model

LR	LR SCHEDULER	CONTEXT LEN.	BATCH	TRAINING STEPS
1e-5 $\rightarrow$ 1e-6	Cosine	16,384	128	100

Table 10: NIAH on pure softmax-attention models.

MODEL	METHOD	NIAH-SINGLE-2		NIAH-SINGLE-3	
		32K	64K	32K	64K
Qwen2.5-7B	Pre-train	100.0	95.8	99.8	98.6
	+SFT	100.0	92.6	99.8	99.4
Qwen2.5-3B	Pre-train	100.0	94.6	100.0	99.6
	+SFT	100.0	93.8	100.0	90.8
	+QK-RESTORE	100.0	93.8	100.0	93.4
Mistral-7B	Pre-train	99.8	90.2	99.8	81.0
	+SFT	99.6	59.6	65.8	4.80
	+QK-RESTORE	99.8	62.0	70.0	5.40

(e.g., Qwen2.5-7B drops from 95.8 to 92.6 at 64K), indicating that these models already learn robust long-context retrieval during pre-training. Accordingly, QK-RESTORE yields only marginal gains, consistent with a near-saturation regime where little capability is lost. To investigate the underlying mechanism, we measure the attention gradient decay on Qwen2.5-3B, shown in Figure 7. Unlike HypeNet-2B, whose full-attention layers exhibit a well-fitted exponential decay, it has a clear near-flat tail. We hypothesize that long-context recall in

pure softmax-attention models is maintained structurally, via the abundance of full-attention layers.

N Method of Identifying Retrieval Layers in Pure Softmax-Attention Models

To identify which softmax-attention layers are primarily responsible for long-range retrieval, we conduct a leave-one-out ablation study over the set L_{attn} . For each layer $\ell \in L_{\text{attn}}$, we construct an ablated model $\mathcal{M}^{(-\ell)}$ by setting $\mathbf{W}_Q^\ell = \mathbf{W}_K^\ell = \mathbf{0}$. Each ablated model is evaluated on several retrieval benchmarks (Ruler-QA-SQuAD (Hsieh et al., 2024), FDA (Arora et al., 2023), SWDE (Lockard et al., 2019)). The importance score of layer ℓ is defined as:

$$\delta\text{Recall}(\ell) = \mathcal{S}(\mathcal{M}) - \mathcal{S}(\mathcal{M}^{(-\ell)}),$$

where $\mathcal{S}(\cdot)$ denotes the average retrieval accuracy across benchmarks. A large $\delta\text{Recall}(\ell)$ indicates that layer ℓ contributes disproportionately to retrieval and that its routing weights are critical; a near-zero score indicates redundancy. Layers are ranked by δRecall in descending order, and the top- k layers are selected as the target set for weight restoration. In our experiments, we keep $k = L_{\text{attn}}/4$ as the same as in HypeNet.

O Details of deriving QK-PRO

The bilinear constraint $\mathbf{W}_Q^{\text{new}} \mathbf{W}_K^{\text{new}\top} = \mathbf{R}_{\text{pre}}$ has no closed-form joint solution, so we linearise by

fixing one factor to its pre-SFT value, making the remaining problem a standard constrained least-squares system solvable via Lagrange multipliers.

We fix $\mathbf{W}_K^{\text{new}} = \mathbf{W}_K^{\text{pre}}$, find minimum-deviation $\mathbf{W}_Q^{\text{new}}$:

$$\min_{\mathbf{W}_Q^{\text{new}}} \left\| \mathbf{W}_Q^{\text{new}} - \mathbf{W}_Q^{\text{post}} \right\|_F \quad (91)$$

$$\text{s.t.} \quad \mathbf{W}_Q^{\text{new}} \mathbf{W}_K^{\text{new}\top} = \mathbf{R}_{\text{pre}} \quad (92)$$

Then applying the Lagrange multipliers:

$$\mathcal{L} = \left\| \mathbf{X} - \mathbf{W}_Q^{\text{post}} \right\|_F + \text{tr}(\Lambda^\top (\mathbf{X} \mathbf{W}_K^{\text{new}\top} - \mathbf{R}_{\text{pre}})) \quad (93)$$

Taking the derivative $\partial \mathcal{L} / \partial \mathbf{X} = 0$, and we get:

$$\mathbf{W}_Q^{\text{new}} = \mathbf{W}_Q^{\text{post}} + (\mathbf{R}_{\text{pre}} - \mathbf{W}_Q^{\text{post}} \mathbf{W}_K^{\text{pre}\top}) (\mathbf{W}_K^{\text{pre}} \mathbf{W}_K^{\text{pre}\top} + \lambda \mathbf{I})^{-1} \mathbf{W}_K^{\text{pre}} \quad (94)$$