

Dynamic Long Context Reasoning over Compressed Memory via End-to-End Reinforcement Learning

Zhuoen Chen^{1,†}, Dongfang Li^{1,†}, Meishan Zhang¹, Baotian Hu^{1,*}, Min Zhang¹

¹Harbin Institute of Technology, Shenzhen

Abstract

Large Language Models (LLMs) face severe challenges in long-context processing, including quadratic computational costs, information forgetting, and the context fragmentation inherent in Retrieval-Augmented Generation (RAG). We introduce LycheeMemory, a cognitively inspired framework that enables efficient long-context inference via chunk-wise compression and selective memory recall, rather than processing all raw tokens. LycheeMemory segments the input into chunks and encodes each into compressed KV-cache-style representations using a Compressor. A Gate then dynamically selects relevant memory blocks, which a Reasoner iteratively processes with an evolving working memory to solve downstream tasks. The Compressor and Reasoner are jointly optimized via end-to-end reinforcement learning, while the Gate is trained separately as a classifier. Experimental results demonstrate that LycheeMemory achieves competitive accuracy (up to 82% in ablation variants) on multi-hop reasoning benchmarks (e.g., RULER-HQA), successfully extrapolates context length from 7K to 1.75M, and provides a favorable accuracy–efficiency trade-off against strong long-context baselines. Notably, compared to MemAgent, LycheeMemory achieves an average $2\times$ reduction in peak GPU memory usage and a $6\times$ speedup during inference.

1 Introduction

Despite the remarkable capabilities demonstrated by Large Language Models (LLMs), efficiently processing long contexts remains a critical challenge (Liu et al., 2025; Comanici et al., 2025; Wan et al., 2025). To address this bottleneck, current methodologies primarily diverge into three paradigms, each facing inherent trade-offs between

efficiency and capability. Sparse and linear attention mechanisms (Beltagy et al., 2020; Xiao et al., 2023; Katharopoulos et al., 2020) reduce computational complexity but often suffer from performance degradation on extremely long sequences. Retrieval-Augmented Generation (RAG) (Lewis et al., 2020; Karpukhin et al., 2020; Izacard and Grave, 2021) mitigates length constraints while facing severe context fragmentation. By treating text chunks as independent entities, it disrupts the logical dependencies essential for multi-hop reasoning and struggles to capture implicit semantic connections. Conversely, recurrent architectures like RecurrentGPT (Zhou et al., 2023) and MemAgent (Yu et al., 2025) rely on sequential state updates, resulting in slow serial inference speeds that significantly hinder scalability.

To overcome these limitations, we draw inspiration from the mechanisms of human memory and propose LycheeMemory. By mimicking the division of labor between compressed memory bank (i.e., *long-term memory*) and dynamic working memory (Atkinson and Shiffrin, 1968), our framework splits the input text into chunks and compresses them into efficient, high-fidelity compressed KV-cache representations. This builds a compressed memory bank that preserves semantic information while reducing computational costs. During inference, we use a dynamic recall and reasoning workflow driven by a Gate and a Reasoner. It starts with an empty *working memory*, explicitly instantiated as a fixed-length, token-level context window. This design preserves a discrete action space, thereby enabling the memory update process to be optimized via Reinforcement Learning (RL). Subsequently, LycheeMemory sequentially traverses the compressed memory bank: for each chunk, the Gate evaluates whether the chunk contributes to the current reasoning state, given the current working memory and the user query. If deemed relevant, the Reasoner utilizes the chunk

[†] Equal contribution. *Correspondence: hubaotian@hit.edu.cn. Code available at <https://github.com/owoakuma/LycheeMemory>.

to update current working memory; otherwise, the chunk is skipped. Through this selective update and iterative refinement, the Reasoner facilitates multi-step reasoning across multiple memory chunks, avoiding the blind processing of the entire input sequence characteristic of traditional recurrent architectures.

A core challenge is ensuring that the compressed memory can be effectively used by the Reasoner for precise inference. We adopt a joint policy optimization strategy: we train the Compressor and Reasoner end-to-end with RL, and train the Gate separately as a classifier. We evaluate LycheeMemory on RULER-HQA (Yang et al., 2018; Hsieh et al., 2024), 2WikiMultihopQA (Ho et al., 2020), and StreamingQA (Liska et al., 2022). Experimental results show that LycheeMemory maintains competitive accuracy on multi-hop reasoning, extrapolates context length from 7K to 1.75M, and improves the accuracy–efficiency trade-off. Compared to MemAgent (Yu et al., 2025), LycheeMemory reduces peak GPU memory usage by 2× and speeds up inference by 6×.

The main contributions of this work are summarized as follows:

- We propose LycheeMemory, a framework comprising a Compressor, a Gate, and a Reasoner, which transforms long-context processing from direct modeling of raw tokens into efficient iterative reasoning over a compressed memory bank.
- We introduce a joint policy optimization strategy that trains the Compressor and Reasoner end-to-end via RL, enabling the compressed memory to be directly optimized for downstream tasks.
- Experimental results show that LycheeMemory scales the context size to 1.75M tokens and improves inference efficiency while maintaining competitive accuracy.

2 Related Work

Explicit Memory Methods. Explicit memory methods externalize context as human-readable text or symbols. Standard RAG retrieves static chunks via semantic similarity but often suffers from context fragmentation and limited precision in multi-hop reasoning (Gutiérrez et al., 2025; Weller et al., 2025; Merola and Singh, 2025). Agentic

memory systems mitigate this by actively managing external memory, such as MemGPT’s OS-inspired hierarchy (Packer et al., 2023) and Mem0’s lifecycle-based memory updates (Chhikara et al., 2025). More recent RL-based approaches (e.g., MemAgent (Yu et al., 2025), Mem1 (Zhou et al., 2025)) learn to manage a bounded memory by selectively overwriting or integrating observations during streaming. Despite their interpretability, these methods operate on raw tokens and incur substantial computational overhead. In contrast, our approach leverages compressed memory with selective retrieval, achieving lower peak memory usage and inference latency.

Implicit Memory Methods. Implicit memory methods optimize internal representations via activation compression or parametric updates. To alleviate the quadratic cost of self-attention, cache compression approaches exploit attention sparsity, retaining only salient tokens (e.g., H2O (Zhang et al., 2023), SnapKV (Li et al., 2024)). Beyond static pruning, dynamic methods retrieve relevant cache blocks on demand (Xiao et al., 2024; Gao et al., 2025). Parametric alternatives, such as DyPRAG (Tan et al., 2025), encode documents into latent LoRA adapters (Hu et al., 2022) and route queries to specialized weights. While effective in reducing memory footprint, aggressive compression often degrades long-tail reasoning (Zhang et al., 2025), and purely latent approaches (Hao et al., 2024; Eyuboglu et al., 2025) lack structured retrieval needed for large-scale multi-document streams. In contrast, our method couples selective retrieval with iterative working memory updates via a Gate and Reasoner, enabling robust multi-hop reasoning over million-token contexts.

Overall, LycheeMemory bridges explicit and implicit memory: it stores documents as compressed KV-cache representations, while performing state-dependent retrieval and reasoning through a plain-text working memory. This retains the scalability benefits of compression and yields an interpretable trace over selected evidence chunks.

3 Methodology

3.1 Overview

We address long-context modeling where a model takes ultra-long documents D (length N) and a user query Q to generate an answer A . Due to the prohibitive length of D , processing the entire sequence directly is computationally infeasible. To

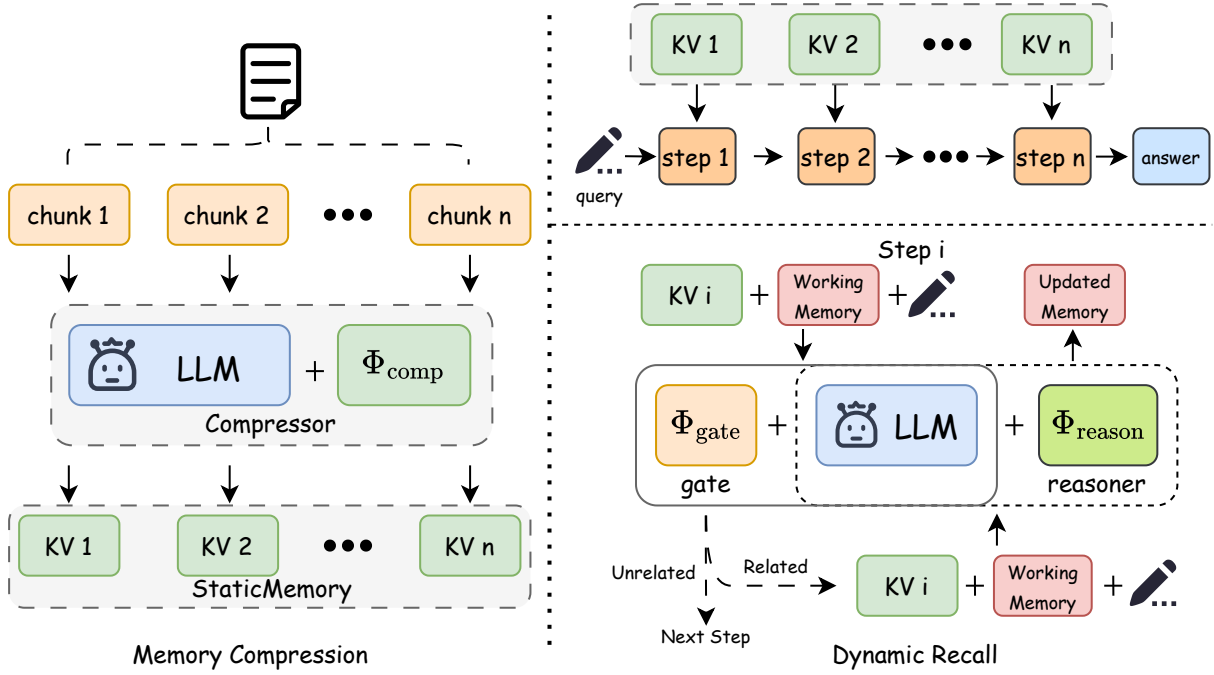


Figure 1: Overview of the LycheeMemory framework. The left panel illustrates compressed memory construction, where a long document is segmented and compressed into compact KV-cache representations by the compressor. The right panel depicts the dynamic recall and reasoning workflow, in which the gate selectively activates relevant memory blocks and the reasoner iteratively updates the working memory to produce the final answer.

address this, we propose **LycheeMemory**, a dual-system framework for long-context processing. As illustrated in Figure 1, the architecture comprises three core roles:

- Compressor Φ_{comp} : Composed of the base model Φ augmented with a compression LoRA module Ψ_{comp} , responsible for encoding raw text into KV-cache-style memory.
- Gate Φ_{gate} : Composed of the base model Φ augmented with a gating LoRA module Ψ_{gate} , acting as a relevance filter.
- Reasoner Φ_{reason} : Composed of the base model Φ augmented with a reasoning LoRA module Ψ_{reason} , responsible for complex reasoning based on recalled memories.

Let D be segmented into K sequential chunks (size sz , i.e., $K = N/sz$) as $D = \{C_1, C_2, \dots, C_K\}$. In our experiments, we set $sz = 4096$. The processing workflow of LycheeMemory involves two main phases:

Memory Compression: In this phase (detailed in §3.2), each text chunk C_k is processed by the Compressor Φ_{comp} and encoded into a compact latent representation θ_k . This representation is subsequently stored in the compressed memory bank Θ , i.e., $\Theta = \{\theta_1, \dots, \theta_K\}$.

Dynamic Recall and Reasoning: Distinct from the latent representations used for storage, the model maintains a working memory \mathbf{m} during the dynamic recall and reasoning phase. \mathbf{m} exists as plaintext tokens within the model’s context and is iteratively updated as the model scans the compressed memory bank to maximize reasoning capability. When receiving a user query Q (detailed in §3.3), the model scans memory blocks with index $i = 1, \dots, K$, activating the Gate Φ_{gate} and Reasoner Φ_{reason} . The process starts with an initial empty working memory \mathbf{m}_0 . At scan step i , the Gate evaluates the compressed memory block θ_i in conjunction with the current working memory \mathbf{m}_t and query Q . If deemed relevant, the Reasoner is invoked to update the working memory state: $\mathbf{m}_{t+1} = \Phi_{\text{reason}}(\mathbf{m}_t, \theta_i, Q)$, and we increment the update index $t \leftarrow t + 1$; otherwise, we skip this block and keep t unchanged. Finally, the model synthesizes the answer A based on \mathbf{m}_T and Q , where $T \leq K$.

3.2 Compressed Memory Construction

The construction of the compressed memory bank Θ is central to the LycheeMemory framework. We present a KV-cache compression style method that achieves an optimal balance between information density and computational efficiency.

3.2.1 KV-cache Style Compression via Memory Tokens

Similar to previous works (Chevalier et al., 2023; Deng et al., 2025), we define the compression as a mapping from text to a latent representation, $C_i \rightarrow \theta_i$. We utilize base model Φ augmented by a LoRA module Ψ_{comp} as the Compressor, eliminating the need for an external encoder. For any text chunk $C_i = [x_1^i, \dots, x_w^i]$ of length w , we first determine a compression ratio α_i . We then define a set of $z_i = w/\alpha_i$ trainable memory tokens $V_i = \{\langle v \rangle_1^i, \dots, \langle v \rangle_{z_i}^i\}$. Next, we interleave V_i with C_i by inserting a memory token after every α_i original tokens, forming interleaved sequence C'_i :

$$\begin{aligned} C'_i &= \text{Interleave}(C_i, V_i) \\ &= [x_1^i, \dots, x_{\alpha_i}^i, \langle v \rangle_1^i, \dots, x_w^i, \langle v \rangle_{z_i}^i] \end{aligned}$$

This sequence C'_i is passed through Φ_{comp} for a single forward pass. During this process, the model is trained to embed the semantic information of the preceding α_i tokens into the hidden state of the subsequent memory token $\langle v \rangle_j^i$. Finally, the set of hidden states corresponding to all memory tokens constitutes the compact KV-cache style representation θ_i stored in the compressed memory bank Θ :

$$\begin{aligned} \theta_i &= \{h(\langle v \rangle_1^i), \dots, h(\langle v \rangle_{z_i}^i)\}, \\ \text{where } h(\cdot) &= \text{HiddenState}(\Phi_{\text{comp}}(C'_i)). \end{aligned}$$

3.2.2 Pre-optimization of the Compressor

Before end-to-end RL, to ensure that θ_i retains the core semantic information of C_i despite high compression, we jointly optimize the LoRA module Ψ_{comp} while keeping the base model Φ frozen using data augmentation and diverse tasks. Note that in the encoding phase ($C'_i \rightarrow \theta_i$), the base model Φ combined with Ψ_{comp} generates the compressed representation θ_i . Conversely, in all subsequent decoding tasks based on θ_i , we utilize only the frozen base model Φ without Ψ_{comp} for generation. This design ensures that the gradient flows only through Ψ_{comp} , effectively decoupling the compression capability from the general generation ability of the base model. Given a compressed representation θ_i , the model Φ is trained to perform three distinct tasks. Let $P_\Phi(Y|\text{context})$ be the probability generating Y given the context:

Text Reconstruction. The model must regenerate the original text C_i using only θ_i as context.

$$\mathcal{L}_{\text{recon}} = -\log P_\Phi(C_i|\theta_i)$$

QA Generation. We pre-generate synthetic question-answer pairs (Q_j, A_j) for C_i . The model generates A_j given θ_i and Q_j . The loss \mathcal{L}_{qa} is computed only over the answer A_j .

$$\mathcal{L}_{\text{qa}} = -\mathbb{E}_{(Q_j, A_j) \sim C_i} [\log P_\Phi(A_j|\theta_i, Q_j)]$$

Creative Generation. The model performs high-level semantic tasks based on θ_i , such as generating a summary S_i . We use the model output based on the original text, $\Phi(C_i)$, as the ground-truth label Y_{creative} .

$$\begin{aligned} \mathcal{L}_{\text{creat}} &= -\log P_\Phi(Y_{\text{creat}}|\theta_i) \\ \text{where } Y_{\text{creat}} &= \Phi(C_i, P_{\text{task}}) \end{aligned}$$

The total loss $\mathcal{L}_{\text{comp}}$ is a weighted sum of the above losses, minimized by updating Ψ_{comp} :

$$\min_{\Psi_{\text{comp}}} \mathcal{L}_{\text{comp}} = \mathbb{E}_{C_i \sim \mathcal{D}} [w_1 \mathcal{L}_{\text{recon}} + w_2 \mathcal{L}_{\text{qa}} + w_3 \mathcal{L}_{\text{creat}}]$$

We train separate projection matrices for the memory tokens v_{mem} , functionally isolating them from regular token representations to learn a dedicated compression subspace.

3.3 Dynamic Recall and Reasoning Workflow

After constructing the compressed memory bank Θ , the core of LycheeMemory lies in efficiently retrieving and reasoning over these compressed representations. In contrast to methods like MemAgent (Yu et al., 2025), which employ linear scanning with forced updates for every chunk, we introduce a relevance threshold τ . As the system traverses the compressed memory bank, the Gate scores each compressed memory block, and only blocks exceeding this threshold trigger the Reasoner to update the working memory.

3.3.1 LoRA Gate

To avoid the overhead of unnecessary memory updates, we require a filter to discard static blocks irrelevant to the user query Q . An intuitive solution would be an embedding model calculating cosine similarity between chunks and Q . While such lightweight retrieval can be reasonably strong on recall, it only captures static semantic similarity and lacks state-dependent retrieval conditioned on the evolving working memory \mathbf{m} (See §4). This limitation becomes salient in multi-hop settings, where later-hop evidence may only become relevant after intermediate entities are added into \mathbf{m} . A further critical limitation is that external retrievers

cannot leverage the working memory \mathbf{m} , which often contains key secondary clues (e.g., intermediate entities) derived from the query and previously processed memory chunks. Motivated by this, we implement the Gate Φ_{gate} by adding a LoRA adapter Ψ_{gate} to the base model Φ .

Architecture and Inference. Given a user query Q , the current working memory \mathbf{m}_t , and a candidate memory block $\theta_i \in \Theta$ (represented by its memory tokens), we concatenate them and extract the hidden state of the final token, \mathbf{h}_{last} . This state is projected by a trainable linear head \mathbf{W}_{gate} followed by a sigmoid activation to produce a relevance probability:

$$P = \sigma(\mathbf{W}_{\text{gate}} \cdot \mathbf{h}_{\text{last}}(\Phi(Q, \mathbf{m}_t, \theta_i; \Psi_{\text{gate}})))$$

The memory block is used to update the working memory only if $P > \tau$.

Training Objective. Due to the gradient discontinuity caused by discrete recall decisions, we treat gating training as a separate binary classification task beyond RL in §3.4. We align text chunks with downstream tasks (e.g., QA pairs) to construct training data. A memory block θ_i is labeled positive ($y_i^* = 1$) if it contains evidence required to answer Q , and negative ($y_i^* = 0$) otherwise. We optimize the gate parameters (LoRA Ψ_{gate} and Head \mathbf{W}_{gate}) using Binary Cross-Entropy (BCE) loss:

$$\mathcal{L}_{\text{gate}} = -\frac{1}{N} \sum_{i=1}^N \left[y_i^* \log P_i + (1 - y_i^*) \log(1 - P_i) \right]$$

where P_i is the predicted probability. This lightweight design ensures the model identifies memory chunks relevant to the query and current reasoning state in the latent space.

3.4 End-to-End RL Optimization

To empower LycheeMemory with the capability of complex reasoning over compressed memories, we propose an enhanced reinforcement learning framework. Unlike prior approaches that optimize components in isolation, we formulate the entire lifecycle from memory compression to reasoning as a unified joint policy optimization problem. This allows the gradient from the final reasoning outcome to backpropagate through the recall workflow and update the Compressor, ensuring the Θ (i.e., *long-term memory*) is optimized specifically for downstream inference.

3.4.1 Joint Policy Formulation

We define the joint policy π_{ϑ} parameterized by ϑ , which encompasses both the Compressor parameters (Ψ_{comp}) and the Reasoner parameters (Ψ_{reason}). For a given input document D and query Q , the generation of an answer A involves a hierarchical trajectory:

$$\pi_{\vartheta}(A, \mathcal{M}, \Theta \mid D, Q) = \underbrace{\prod_{k=1}^K \pi_{\text{comp}}(\theta_k \mid C_k)}_{\text{Memory Construction}} \cdot \underbrace{\prod_{t=1}^T \pi_{\text{reason}}(\mathbf{m}_t \mid \mathbf{m}_{t-1}, \Theta, Q)}_{\text{Dynamic Recall and Reasoning}}$$

where $\Theta = \{\theta_k\}$ represents the compressed memory bank, and $\mathcal{M} = \{\mathbf{m}_t\}_{t=0}^T$ represents the sequence of working memory updates. Our goal is to maximize the expected reward of the final answer A by optimizing ϑ .

The Unified Objective Function. We formulate the unified objective to jointly optimize compression and reasoning:

$$\mathcal{J}(\vartheta) = \mathbb{E}_{Q \sim \mathcal{D}, \{O_i\}_{i=1}^G \sim \pi_{\vartheta_{\text{old}}}} \left[\frac{1}{G} \sum_{i=1}^G \frac{1}{n_i} \sum_{j=1}^{n_i} (\mathcal{L}_{i,j}^{\text{CLIP}}(\vartheta) - \beta D_{\text{KL}}(\pi_{\vartheta} \parallel \pi_{\text{ref}})) \right]$$

where $\mathcal{L}_{i,j}^{\text{CLIP}}(\vartheta) = \min(\rho_{i,j}(\vartheta) \hat{A}_{i,j}, \text{clip}(\rho_{i,j}(\vartheta), 1 - \epsilon, 1 + \epsilon) \hat{A}_{i,j})$

Here, $\rho_{i,j}(\vartheta)$ represents the sequence-level importance sampling weight defined by GSPO (Zheng et al., 2025). G denotes the group size (number of sampled trajectories per prompt), and n_i denotes the number of tokens in the i -th trajectory. By maximizing $\mathcal{J}(\vartheta)$, the model learns to compress context into Θ such that the reasoning policy maximizes the likelihood of high-advantage trajectories.

3.5 Complexity and Efficiency Analysis

In this section, we analyze the computational efficiency of LycheeMemory compared to existing long-context methods.

Memory Construction. This phase incurs $O(N)$ complexity, but it is a one-time, fully parallelizable pre-processing cost.

Gate. While approaches like MemAgent (Yu et al., 2025) achieve $O(N/sz)$ linear complexity via streaming, they require performing full token generation (i.e., memory updates) for every text chunk. In contrast, although the Gate in LycheeMemory must also traverse all K compressed memory blocks to determine relevance, maintaining an $O(N/sz)$ complexity, the computational cost per block is drastically reduced. The Gate requires only a single forward pass for scalar classification, rather than the computationally expensive autoregressive generation used in standard streaming methods.

Dynamic Recall and Reasoning. The heavy computational load of the Reasoner is decoupled from the document length N and depends only on the number of retrieved blocks T :

$$\mathcal{O}_{\text{inference}} \approx O(N/sz \times C_{\text{gate}} + T \times C_{\text{reason}})$$

where $C_{\text{gate}} \ll C_{\text{reason}}$. Since the Gate efficiently filters out irrelevant information ($T \ll N/sz$), LycheeMemory achieves a significantly lower constant factor in its linear scaling compared to methods that reason over every chunk.

4 Experiments

In this section, we evaluate LycheeMemory on long-context QA, analyze inference efficiency and zero-shot generalization, and validate core design choices through ablations.

4.1 Experimental Setup

Model Configuration We use Qwen2.5-Instruct (Yang et al., 2025b) as the base model and train LycheeMemory-3B/LycheeMemory-7B initialized from Qwen2.5-3B/7B-Instruct.

Dataset Construction Following MemAgent (Yu et al., 2025), we synthesize long-document training data from RULER-HQA (Yang et al., 2018; Hsieh et al., 2024) by mixing query-relevant articles with distractors (avg. 20K tokens). We evaluated contexts from 7K to 1.75M tokens for length extrapolation and reported zero-shot results on 2WikiMultihopQA (Ho et al., 2020), StreamingQA (Liska et al., 2022).

Baselines We compare with Search-R1 (Jin et al., 2025), MemAgent (Yu et al., 2025), DeepSeek-R1-Distill-Qwen (Guo et al., 2025), Qwen-2.5-Instruct-1M (Yang et al., 2025a), and QwenLong-L1 (Wan

et al., 2025), using official configurations. Additional details are in Appendix A and Appendix A.4.

4.2 Main Results

We first evaluate LycheeMemory on the synthesized HotpotQA dataset as context length grows. Table 1 shows the comparison with baselines.

Performance at Scale We compare models from 7K to 896K context lengths. For memory-based models (Search-R1, MemAgent, and LycheeMemory), we further evaluate extrapolation at an ultra-long 1.75M tokens to inspect generalization beyond standard training ranges. As shown in Table 1, several baselines fail even within their nominal windows. Reasoning models (e.g., DS-Distill-Qwen series) degrade rapidly as context length increases. In contrast, MemAgent and LycheeMemory show strong length extrapolation, with only mild performance drop as input length increases, validating the effectiveness of the chunked memory mechanism.

Comparison with MemAgent Compared to MemAgent, our LycheeMemory-7B w/o Gate ablation achieves higher accuracy across most evaluated context lengths, while LycheeMemory with Gate trades a small accuracy drop for substantially improved inference efficiency (see §4.3). This indicates that compressed memory with RL-trained reasoning is competitive in accuracy, and the Gate provides an effective accuracy–efficiency trade-off in ultra-long contexts.

4.3 Inference Efficiency Analysis

A key advantage of LycheeMemory is computational efficiency. We measure end-to-end inference time on $2 \times$ A100 (80GB) for 128 samples from 8K to 128K tokens (generation length 1024, largest non-OOM batch). The reported time includes compression and I/O. Figure 2 shows three regimes:

Quadratic Explosion The Qwen2.5-7B baseline exhibits the expected $O(N^2)$ latency growth. At 64K it is markedly slower than memory-based methods and at 128K it further fails due to OOM.

Linear Growth MemAgent and our ablation LycheeMemory without Gate (linear scan over all compressed memory blocks) show linear $O(N)$ complexity. However, LycheeMemory without Gate remains faster than MemAgent because our compressed memory is a highly compressed KV-cache ($\alpha \gg 1$), so the effective sequence length

Model	7K	14K	28K	56K	112K	224K	448K	896K	1.75M
QwenLong-L1-32B (Wan et al., 2025)	72.66	75.00	72.66	60.94	31.25	17.19	13.28	11.72	OOM
QwenLong-L1.5-30B-A3B (Shen et al., 2025)	67.97	68.75	66.41	63.28	61.72	59.38	43.75	39.06	28.13
Qwen2.5-Instruct-14B-1M (Yang et al., 2025a)	60.16	60.94	50.00	57.03	50.00	37.50	8.59	0.00	OOM
Qwen2.5-Instruct-7B-1M (Yang et al., 2025a)	61.72	56.25	53.91	55.47	51.56	33.59	12.50	0.00	OOM
DS-Distill-Qwen-32B (Guo et al., 2025)	70.31	66.41	65.62	46.88	23.44	13.28	7.81	7.03	OOM
DS-Distill-Qwen-14B (Guo et al., 2025)	64.06	64.84	57.03	40.62	14.84	8.59	3.12	6.25	OOM
DS-Distill-Qwen-7B (Guo et al., 2025)	30.47	12.50	3.12	0.00	0.00	0.78	0.00	0.00	OOM
RAG + Qwen2.5-7B-Instruct	67.19	66.41	66.41	67.19	64.84	64.06	62.5	61.72	62.38
Search-R1 (Jin et al., 2025)	72.66	71.88	67.71	73.96	66.67	62.5	64.58	67.71	67.19
RL-MemAgent-7B (Yu et al., 2025)	82.03	79.69	78.91	77.34	79.69	72.66	74.22	76.56	75.78
LycheeMemory-7B (ours)	77.34 _{1.0x}	76.56 _{1.2x}	75.00 _{1.6x}	76.56 _{2.5x}	75.78 _{3.5x}	73.44 _{5.9x}	74.22 _{9.7x}	72.66 _{17.7x}	71.09 _{28.2x}
LycheeMemory-7B w/o Gate (ours)	80.47	81.25	82.03	81.25	80.47	79.69	75.00	75.78	78.12

Table 1: Comparison of main experimental results under different context lengths. All values are normalized sub-EM accuracy (%). Blue indicates the inference speedup of LycheeMemory relative to its without Gate ablation.




Figure 2: Inference latency as context length increases. LycheeMemory exhibits a nearly flat latency curve, in contrast to the quadratic and linear increases observed in the full-context and MemAgent baselines respectively.

processed by the reasoning workflow is much shorter than the text stream of MemAgent.

Near-Constant Inference With the Gate module, LycheeMemory shows striking efficiency. As context grows from 8K to 128K, inference time rises only slightly. Compression and Gate overhead grows linearly (with a tiny coefficient), while the costly reasoning (with memory update) steps run on only a few retrieved blocks. In terms of results, at 128K we achieve a 6 \times speedup over MemAgent and a 3.5 \times speedup over the w/o Gate baseline; meanwhile, Table 1 shows that the accuracy drop on the closest reported bucket (112K) is only 6%. Additional analyses are in Appendix A.5 and Appendix D.

4.4 Zero-shot Generalization

We evaluate LycheeMemory zero-shot on 2Wiki-MultiHopQA and StreamingQA. Table 2 shows strong performance on unseen multi-document reasoning tasks. Due to space limitations, additional

Method	2WikiMultiHopQA			StreamingQA	
	14K	28K	56K	F1	sub-EM
Qwen2.5-Instruct-7B	57.0	42.2	37.5	30.5	23.4
RAG	68.8	64.1	59.4	84.3	67.2
MemAgent	74.2	73.4	71.1	77.9	60.2
LycheeMemory	75.0	70.3	73.4	80.8	73.4

Table 2: Zero-shot comparison results of 2WikiMultiHopQA and StreamingQA.




Figure 3: QA Accuracy across varying context lengths under different compression ratios. The 4 \times ratio (Ours) achieves the optimal balance, matching the stability of 2 \times while significantly outperforming aggressive compression (16 \times).

OOD evaluations of LongBench benchmark (Bai et al., 2024) on Appendix E.

4.5 Ablation Study

To analyze the contribution of each component, we conduct a series of ablation studies using the LycheeMemory-3B model.

4.5.1 Different Compression Ratios

We study the effect of compression ratios ($\alpha \in \{2, 4, 8, 16\}$) on reasoning accuracy over context lengths from 2K to 128K tokens (Figure 3). Results reveal a clear trade-off between memory efficiency and information retention. Both 2 \times and 4 \times

Method	56K	112K	224K
Text-embedding-3-large	94.3	82.1	80.9
Gate (Query Only)	88.2	76.4	74.8
Gate (Query + Memory)	98.5	86.3	84.1

Table 3: Recall of gold supporting chunks on multi-hop QA samples across context lengths. All methods retrieve the top 8 chunks under an identical retrieval budget.

compression maintain near-lossless performance, preserving $> 80\%$ accuracy even at 128K, with a negligible gap ($< 1\%$) between them, indicating that $4\times$ compression is sufficient to capture semantic density without redundancy. In contrast, $16\times$ compression degrades sharply (71.5% at 2K to 42.0% at 128K), while $8\times$ provides a compromise but exhibits mild attrition ($< 10\%$) at extreme lengths. Accordingly, we adopt $\alpha = 4$ as the default, halving the memory footprint of $\alpha = 2$ with no statistically significant loss in reasoning performance.

4.5.2 Ablation on Gate

Experimental Setup. We evaluate different retrieval strategies under increasing context lengths by segmenting the input into non-overlapping 4096-token chunks. For the embedding baseline, we further split each 4096-token chunk into 1024-token micro-chunks, score each micro-chunk with the query, and use the maximum score as the chunk score.

Results and Analysis. As shown in Table 3, all methods perform well at shorter contexts (56K). However, baselines show a clear performance drop as context length increases. Static embedding-based retrieval and query-only Gate decline at 112K, with the strongest baseline dropping to 82.1%. In contrast, our Gate conditioned on both the query and the evolving working memory maintains a high recall of 86.3% at 112K and 84.1% at 224K, consistently surpassing other retrieval strategies.

This trend reflects the state-dependent nature of multi-hop reasoning: static retrievers model $P(\text{Chunk} \mid Q)$ and overemphasize early-hop evidence, whereas LycheeMemory conditions retrieval on the evolving memory state, modeling $P(\text{Chunk} \mid Q, \mathbf{m}_t)$, which enables adaptive evidence discovery across reasoning steps.

Models / Stages	Evaluation Metrics (sub-EM)		
	HotpotQA	2Wiki	Avg.
Qwen2.5-3B-Instruct			
└ Stage-1: Naive Chunking	38.28	35.16	36.72
└ Stage-2: Memory Compression	39.84	36.72	38.28
Stage-3: w/ SFT	60.16	58.59	59.38
Stage-3: w/ RL	68.75	64.84	66.80
Stage-3: w/ End-to-End RL	70.31	67.19	68.75

Table 4: Ablation study of the staged optimization process. The base model is Qwen2.5-3B-Instruct with a fixed context length of 16k tokens.

4.5.3 Analysis of Staged Optimization Strategies

Table 4 analyzes the impact of each training stage. Memory Compression (Stage-2) achieves performance comparable to Naive Chunking (Stage-1) with reduced token usage, indicating that compression alone requires further alignment. Stage-3 SFT yields a notable improvement (+21.10 sub-EM) by learning basic interaction patterns, but is surpassed by RL Optimization, which better supports multi-hop navigation and error correction. The best performance (68.75 Avg. sub-EM) is obtained with End-to-End RL, where joint optimization enables gradients to reach the compressor, encouraging reasoning-aware representations and validating the need for unified perception–reasoning training. We further provide a training convergence analysis for the joint optimization stage in Appendix B.

5 Conclusion

We introduce LycheeMemory, a cognitively inspired framework that enables efficient long-context reasoning by mimicking the human memory’s division into long-term storage and dynamic working memory. Our method integrates a Compressor, a Gate, and a Reasoner: we jointly optimize the Compressor and Reasoner through end-to-end reinforcement learning, and train the Gate separately as a classifier. Experimental results demonstrate that the LycheeMemory w/o Gate ablation can reach up to 82% normalized sub-EM accuracy on multi-hop benchmarks and scales context length to 1.75M tokens, while the full model provides a favorable accuracy–efficiency trade-off. Compared to MemAgent, LycheeMemory provides a $2\times$ reduction in peak GPU memory and a $6\times$ inference speedup. Overall, LycheeMemory offers an efficient solution for ultra-long context modeling.

Limitations

A primary limitation is the trade-off between memory compression and information retention; as shown in our ablations, aggressive compression ratios (e.g., exceeding 16×) can lead to the loss of fine-grained details necessary for complex reasoning. Additionally, while our gating mechanism drastically reduces the computational load of the reasoner, the gate itself must still perform a lightweight linear scan over all memory blocks, maintaining a linear complexity component. Finally, the multi-stage training pipeline, which involves coordinating separate compression and reasoning modules via reinforcement learning, is more complex to optimize than standard supervised fine-tuning approaches.

Ethical Considerations

While LycheeMemory provides advancements in long-context processing efficiency, we acknowledge several ethical considerations associated with its development and deployment. First, the use of memory compression mechanisms introduces a risk of information loss or feature collapse, which could lead to subtle hallucinations or the misrepresentation of facts in sensitive domains such as legal or medical document analysis. This necessitates caution when using the system for high-stakes decision-making. Second, our framework relies on RL to optimize reasoning policies; since RL can inadvertently amplify biases present in the training data or the backbone LLM, practitioners should monitor for biased outputs during deployment. Third, regarding data privacy, the ability of models to process and store extremely long interaction histories increases the potential scope for capturing and retaining sensitive personal information. Users and developers must ensure robust data anonymization and secure storage for the compressed memory banks. Finally, while the multi-stage training pipeline incurs a non-negligible carbon footprint, the significant reduction in inference-time GPU memory and latency contributes to more sustainable and energy-efficient large-scale AI operations in the long run. We encourage the research community to further investigate the interpretability of latent memory representations to ensure transparency and accountability in compressed memory systems.

Acknowledgements

This work was supported by the National Natural Science Foundation of China (Grant No. 62406088, 62422603), Guangdong Basic and Applied Basic Research Foundation (Grant No. 2025A1515011376), Guangdong Basic and Applied Basic Research Foundation (Grant No. 2024B0101050003).

References

- Richard C Atkinson and Richard M Shiffrin. 1968. Human memory: A proposed system and its control processes. In *Psychology of learning and motivation*, volume 2, pages 89–195. Elsevier.
- Yushi Bai, Xin Lv, Jiajie Zhang, Hongchang Lyu, Jiankai Tang, Zhidian Huang, Zhengxiao Du, Xiao Liu, Aohan Zeng, Lei Hou, and 1 others. 2024. Longbench: A bilingual, multitask benchmark for long context understanding. In *Proceedings of the 62nd annual meeting of the association for computational linguistics (volume 1: Long papers)*, pages 3119–3137.
- Iz Beltagy, Matthew E Peters, and Arman Cohan. 2020. Longformer: The long-document transformer. *arXiv preprint arXiv:2004.05150*.
- Alexis Chevalier, Alexander Wettig, Anirudh Ajith, and Danqi Chen. 2023. [Adapting language models to compress contexts](#). In *The 2023 Conference on Empirical Methods in Natural Language Processing*.
- Prateek Chhikara, Dev Khant, Saket Aryan, Taranjeet Singh, and Deshraj Yadav. 2025. [Mem0: Building production-ready ai agents with scalable long-term memory](#). *Preprint*, arXiv:2504.19413.
- Gheorghe Comanici, Eric Bieber, Mike Schaekermann, Ice Pasupat, Noveen Sachdeva, Inderjit Dhillon, Marcel Blistein, Ori Ram, Dan Zhang, Evan Rosen, and 1 others. 2025. Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities. *arXiv preprint arXiv:2507.06261*.
- Chenlong Deng, Zhisong Zhang, Kelong Mao, Shuaiyi Li, Tianqing Fang, Hongming Zhang, Haitao Mi, Dong Yu, and Zhicheng Dou. 2025. [Unigist: Towards general and hardware-aligned sequence-level long context compression](#). In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*.
- Sabri Eyuboglu, Ryan Ehrlich, Simran Arora, Neel Guha, Dylan Zinsley, Emily Liu, Will Tennien, Atri Rudra, James Zou, Azalia Mirhoseini, and 1 others. 2025. Cartridges: Lightweight and general-purpose long context representations via self-study. *arXiv preprint arXiv:2506.06266*.

- Chaochen Gao, Xing W, Qi Fu, and Songlin Hu. 2025. [Quest: Query-centric data synthesis approach for long-context scaling of large language model](#). In *The Thirteenth International Conference on Learning Representations*.
- Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, and 1 others. 2025. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*.
- Bernal Jiménez Gutiérrez, Yiheng Shu, Weijian Qi, Sizhe Zhou, and Yu Su. 2025. [From RAG to memory: Non-parametric continual learning for large language models](#). In *Forty-second International Conference on Machine Learning*.
- Shibo Hao, Sainbayar Sukhbaatar, DiJia Su, Xian Li, Zhiting Hu, Jason Weston, and Yuandong Tian. 2024. [Training large language models to reason in a continuous latent space](#). *CoRR*, abs/2412.06769.
- Xanh Ho, Anh-Khoa Duong Nguyen, Saku Sugawara, and Akiko Aizawa. 2020. Constructing a multi-hop qa dataset for comprehensive evaluation of reasoning steps. *arXiv preprint arXiv:2011.01060*.
- Cheng-Ping Hsieh, Simeng Sun, Samuel Kriman, Shantanu Acharya, Dima Rekesh, Fei Jia, and Boris Ginsburg. 2024. Ruler: What’s the real context size of your long-context language models? In *First Conference on Language Modeling*.
- Edward J Hu, yelong shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2022. [LoRA: Low-rank adaptation of large language models](#). In *International Conference on Learning Representations*.
- Gautier Izacard and Edouard Grave. 2021. Leveraging passage retrieval with generative models for open domain question answering. In *Proceedings of the 16th conference of the european chapter of the association for computational linguistics: main volume*, pages 874–880.
- Bowen Jin, Hansi Zeng, Zhenrui Yue, Jinsung Yoon, Sercan Arik, Dong Wang, Hamed Zamani, and Jiawei Han. 2025. Search-r1: Training llms to reason and leverage search engines with reinforcement learning. *arXiv preprint arXiv:2503.09516*.
- Vladimir Karpukhin, Barlas Oguz, Sewon Min, Patrick SH Lewis, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen-tau Yih. 2020. Dense passage retrieval for open-domain question answering. In *EMNLP (1)*, pages 6769–6781.
- Angelos Katharopoulos, Apoorv Vyas, Nikolaos Pappas, and François Fleuret. 2020. Transformers are rns: Fast autoregressive transformers with linear attention. In *International conference on machine learning*, pages 5156–5165. PMLR.
- Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, and 1 others. 2020. Retrieval-augmented generation for knowledge-intensive nlp tasks. *Advances in neural information processing systems*, 33:9459–9474.
- Yuhong Li, Yingbing Huang, Bowen Yang, Bharat Venkitesh, Acyr Locatelli, Hanchen Ye, Tianle Cai, Patrick Lewis, and Deming Chen. 2024. [SnapKV: LLM knows what you are looking for before generation](#). In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*.
- Adam Liska, Tomas Kocisky, Elena Gribovskaya, Tayfun Terzi, Eren Sezener, Devang Agrawal, Cyprien De Masson D’Autume, Tim Scholtes, Manzil Zaheer, Susannah Young, and 1 others. 2022. Streamingqa: A benchmark for adaptation to new knowledge over time in question answering models. In *International Conference on Machine Learning*, pages 13604–13622. PMLR.
- Jiaheng Liu, Dawei Zhu, Zhiqi Bai, Yancheng He, Huanxuan Liao, Haoran Que, Zekun Wang, Chenchen Zhang, Ge Zhang, Jiebin Zhang, and 1 others. 2025. A comprehensive survey on long context language modeling. *arXiv preprint arXiv:2503.17407*.
- Carlo Merola and Jaspinder Singh. 2025. [Reconstructing context: Evaluating advanced chunking strategies for retrieval-augmented generation](#). *Preprint*, arXiv:2504.19754.
- Charles Packer, Vivian Fang, Shishir G. Patil, Kevin Lin, Sarah Wooders, and Joseph E. Gonzalez. 2023. [Memgpt: Towards llms as operating systems](#). *CoRR*, abs/2310.08560.
- Weizhou Shen, Ziyi Yang, Chenliang Li, Zhiyuan Lu, Miao Peng, Huashan Sun, Yingcheng Shi, Shengyi Liao, Shaopeng Lai, Bo Zhang, and 1 others. 2025. Qwenlong-11. 5: Post-training recipe for long-context reasoning and memory management. *arXiv preprint arXiv:2512.12967*.
- Yuqiao Tan, Shizhu He, Huanxuan Liao, Jun Zhao, and Kang Liu. 2025. [Dynamic parametric retrieval augmented generation for test-time knowledge enhancement](#). *Preprint*, arXiv:2503.23895.
- Fanqi Wan, Weizhou Shen, Shengyi Liao, Yingcheng Shi, Chenliang Li, Ziyi Yang, Ji Zhang, Fei Huang, Jingren Zhou, and Ming Yan. 2025. Qwenlong-11: Towards long-context large reasoning models with reinforcement learning. *arXiv preprint arXiv:2505.17667*.
- Maurice Weber, Dan Fu, Quentin Anthony, Yonatan Oren, Shane Adams, Anton Alexandrov, Xiaozhong Lyu, Huu Nguyen, Xiaozhe Yao, Virginia Adams, and 1 others. 2024. Redpajama: an open dataset for training large language models. *Advances in neural information processing systems*, 37:116462–116492.

- Orion Weller, Michael Boratko, Iftekhar Naim, and Jinhyuk Lee. 2025. [On the theoretical limitations of embedding-based retrieval](#). *Preprint*, arXiv:2508.21038.
- Chaojun Xiao, Pengle Zhang, Xu Han, Guangxuan Xiao, Yankai Lin, Zhengyan Zhang, Zhiyuan Liu, and Maosong Sun. 2024. [InfLLM: Training-free long-context extrapolation for LLMs with an efficient context memory](#). In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*.
- Guangxuan Xiao, Yuandong Tian, Beidi Chen, Song Han, and Mike Lewis. 2023. Efficient streaming language models with attention sinks. *arXiv preprint arXiv:2309.17453*.
- An Yang, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoyan Huang, Jiandong Jiang, Jianhong Tu, Jianwei Zhang, Jingren Zhou, and 1 others. 2025a. Qwen2. 5-1m technical report. *arXiv preprint arXiv:2501.15383*.
- 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, Jiayi Yang, Jingren Zhou, Junyang Lin, Kai Dang, and 23 others. 2025b. [Qwen2.5 technical report](#). *Preprint*, arXiv:2412.15115.
- Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William Cohen, Ruslan Salakhutdinov, and Christopher D Manning. 2018. Hotpotqa: A dataset for diverse, explainable multi-hop question answering. In *Proceedings of the 2018 conference on empirical methods in natural language processing*, pages 2369–2380.
- Hongli Yu, Tinghong Chen, Jiangtao Feng, Jiangjie Chen, Weinan Dai, Qiyang Yu, Ya-Qin Zhang, Wei-Ying Ma, Jingjing Liu, Mingxuan Wang, and 1 others. 2025. Memagent: Reshaping long-context llm with multi-conv rl-based memory agent. *arXiv preprint arXiv:2507.02259*.
- Peitian Zhang, Zheng Liu, Shitao Xiao, Ninglu Shao, Qiwei Ye, and Zhicheng Dou. 2025. [Long context compression with activation beacon](#). In *The Thirteenth International Conference on Learning Representations*.
- Zhenyu Zhang, Ying Sheng, Tianyi Zhou, Tianlong Chen, Lianmin Zheng, Ruisi Cai, Zhao Song, Yuandong Tian, Christopher Re, Clark Barrett, Zhangyang Wang, and Beidi Chen. 2023. [H2o: Heavy-hitter oracle for efficient generative inference of large language models](#). In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Chujie Zheng, Shixuan Liu, Mingze Li, Xiong-Hui Chen, Bowen Yu, Chang Gao, Kai Dang, Yuqiong Liu, Rui Men, An Yang, and 1 others. 2025. Group sequence policy optimization. *arXiv preprint arXiv:2507.18071*.
- Wangchunshu Zhou, Yuchen Eleanor Jiang, Peng Cui, Tiannan Wang, Zhenxin Xiao, Yifan Hou, Ryan Cotterell, and Mrinmaya Sachan. 2023. Recurrentgpt: Interactive generation of (arbitrarily) long text. *arXiv preprint arXiv:2305.13304*.
- Zijian Zhou, Ao Qu, Zhaoxuan Wu, Sunghwan Kim, Alok Prakash, Daniela Rus, Jinhua Zhao, Bryan Kian Hsiang Low, and Paul Pu Liang. 2025. [Mem1: Learning to synergize memory and reasoning for efficient long-horizon agents](#). *Preprint*, arXiv:2506.15841.

A Implementation Details

This appendix provides the technical specifications necessary for reproducing LycheeMemory. We detail the three-stage training pipeline: (1) Pre-training of the Compressor with **synthetic supervision** (QA pairs generated via self-annotation), (2) Joint Reinforcement Learning of the Compressor and Reasoner, and (3) **supervised** training of the Gate as a binary classifier.

All models are initialized from the Qwen2.5-3B-Instruct and Qwen2.5-7B-Instruct base models. We utilize $2 \times$ NVIDIA A100 (80GB) GPUs for training.

A.1 Stage 1: Compressor Pre-training

The objective of the Compressor is to encode textual information into the latent space of memory tokens. We employ a random compression ratio $\alpha \in \{2, 4, 8, 16\}$, meaning one memory token is inserted for every α text tokens.

Data Construction. We sample up to 1B tokens from the RedPajama (Weber et al., 2024), and train on approximately 160M effective tokens. For each document, we create splits of sizes 2048, 4096, and 8192, denoted as C_i . We use self-annotation to generate synthetic Question-Answer pairs (i.e., synthetic supervision), serving as training targets for reconstruction and comprehension tasks.

Configuration. We train a LoRA adapter (Ψ_{comp}) for the Compressor with a rank of $r = 64$ and a LoRA alpha of $\alpha_{\text{lora}} = 128$. The larger rank is selected to ensure sufficient representation capacity for the compression task. We use the AdamW optimizer with an initial learning rate of $5e-5$ and a cosine annealing schedule, with the maximum learning rate set to $1e-4$. The batch size is set to 8, and training proceeds for 5,000 steps.

A.2 Stage 2: Joint Reinforcement Learning Optimization

We employ the GSPO algorithm for training. The full procedure is summarized in Algorithm 1. The chunk size is set to 4096, the rollout batch size to 128, the group size G to 12, and the update batch size to 16. The KL divergence coefficient β is set to $1e-3$. We use the AdamW optimizer. Since LoRA is the optimization target, we set the learning rate to $3e-5$ with a linear warmup scheduler over 10 steps. In our runs, the joint optimization typically converges within 150 optimizer update steps and

takes about three days of wall-clock time on $2 \times$ A100 (80GB).

Reward Configuration. During training, we employ a strict rule-based reward validator to prevent reward hacking. We extract tokens within the `<answer></answer>` tags of the final output. If the extracted answer matches the ground truth exactly, every update step in the trajectory receives a reward of 1.0; otherwise, the reward is 0.0. We adopt this stricter validator during RL to avoid exploiting normalization artifacts that are acceptable for evaluation.

Dataset Construction. We follow the dataset construction methodology of MemAgent. Each training sample consists of 130 documents from HotpotQA, with a total token length of approximately 20K. We thoroughly cleaned the dataset by filtering out questions where Qwen2.5-7B-Instruct could achieve a 100% Best-Of-2 score without any context (zero-shot). We selected the top 32,768 processed samples as our training set. Similarly, we synthesized 192 samples from the HotpotQA validation set. For extrapolation testing, we used the same pipeline to synthesize test sets with varying context lengths, where the number of Wikipedia entries ranges from 50 to 3,200, corresponding to context lengths from approximately 7K to 1.75 million tokens.

A.3 Stage 3: Gate Module Training

The Gate is trained separately as a binary classifier to determine whether a memory block has retrieval and reasoning value given the current query and working memory.

Label Assignment. Training data is derived from the rollout process in the RL stage. For multi-hop questions, chunk updates (i.e., memory block updates) containing supporting facts are labeled as **Positive** ($y = 1$). Chunk updates containing no supporting facts are labeled as **Negative** ($y = 0$).

Objective. We minimize the Binary Cross-Entropy (BCE) loss. To mitigate the class imbalance problem (where irrelevant paragraphs far outnumber relevant ones), we apply a positive class weight of $pos_weight = 3.0$.

Configuration. The Gate LoRA adapter (Ψ_{gate}) uses a smaller rank of $r = 16$. We train for 3 epochs with a learning rate of $5e-5$. During inference, the gating threshold τ is empirically set to

Algorithm 1 Joint Policy Optimization for LycheeMemory (GSPO)

Require: Joint policy $\pi_{\vartheta} = \{\pi_{\text{comp}}, \pi_{\text{reason}}\}$, reference model π_{ref} (frozen), dataset \mathcal{D} , group size G , clipping ϵ , KL coefficient β

Ensure: Optimized parameters ϑ

```
1: while not converged do
2:   Sample document–query pair  $(D, Q) \sim \mathcal{D}$ 
3:   for  $g = 1$  to  $G$  do                                     ▷ Group sampling for the same  $(D, Q)$ 
4:     Sample memory  $\Theta_g \sim \pi_{\text{comp}}(\cdot | D)$ 
5:     Sample answer  $A_g \sim \pi_{\text{reason}}(\cdot | \Theta_g, Q)$ 
6:     Define trajectory  $y_g = (\Theta_g, A_g)$ 
7:     Compute reward  $\hat{r}_g = \mathcal{R}(Q, A_g)$ 
8:     Compute KL penalty  $d_g = D_{\text{KL}}(\pi_{\vartheta}(y_g) \| \pi_{\text{ref}}(y_g))$ 
9:      $r_g \leftarrow \hat{r}_g - \beta d_g$ 
10:   end for
11:    $\{\hat{A}_g\}_{g=1}^G \leftarrow \text{GROUPNORM}(\{r_g\}_{g=1}^G)$ 
12:   for  $g = 1$  to  $G$  do
13:      $\rho_g \leftarrow \frac{\pi_{\vartheta}(y_g)}{\pi_{\vartheta_{\text{old}}}(y_g)}$ 
14:      $\mathcal{J}_g \leftarrow \min(\rho_g \hat{A}_g, \text{clip}(\rho_g, 1 - \epsilon, 1 + \epsilon) \hat{A}_g)$ 
15:   end for
16:    $\vartheta \leftarrow \vartheta + \eta \nabla_{\vartheta} \frac{1}{G} \sum_{g=1}^G \mathcal{J}_g$ 
17: end while
```

0.5.

A.4 Evaluation and Baselines

Evaluation Metrics. During evaluation, we report normalized sub-EM (Exact Match). We normalize both the model answer and the ground truth (e.g., removing definite articles, ignoring case differences) and compute a sub-EM score. This means if an answer contains all elements of the standard answer, it is considered correct. When an answer consists of multiple parts, the score corresponds to the proportion of correct parts provided.

Long-Context Benchmarks. We evaluate our model on three long-context QA benchmarks, including RULER-HQA (Yang et al., 2018; Hsieh et al., 2024), 2WikiMultihopQA (Ho et al., 2020), and StreamingQA (Liska et al., 2022). Below we describe the benchmark construction and our implementation details.

- **RULER-HQA:** A synthetic long-context HotpotQA benchmark derived from the RULER framework. Similar to HotpotQA, each query has two supporting documents (gold evidence). We construct long contexts by mixing the gold evidence with irrelevant distractor documents (sourced from other samples). We build test sets

with varying total context lengths ($N \in \{7k, 14k, 28k, \dots, 448k, 896k, 1.75M\}$), with randomized evidence positions.

- **2WikiMultihopQA:** A multi-hop QA dataset built from Wikipedia. We follow the same long-context construction and evaluation pipeline as RULER-HQA: we take the query-relevant evidence documents from 2WikiMultihopQA and mix them with distractor documents to reach the target context length (14K/28K/56K in our experiments). We use the same chunking setting ($sz = 4096$), memory compression, dynamic recall, and normalized sub-EM evaluation.
- **StreamingQA:** A streaming QA benchmark designed for evaluation under continuously growing corpora. For our long-context setting, we concatenate the documents of all questions into a single global document of approximately 800k tokens, and evaluate each question by running LycheeMemory over this shared 800k context. We use the same chunking setting ($sz = 4096$) and normalized sub-EM evaluation.

Baselines. We compare LycheeMemory against three categories of strong baselines:

- **RAG Agent:** We implement a standard Retrieval-Augmented Generation agent using OpenAI’s `text-embedding-3-large` as the retriever. The document is segmented into semantic chunks (Wikipedia entries). For each query, the agent retrieves the top-8 most relevant chunks and feeds them into the base model for generation.
- **Search-R1:** We use a Search-R1 agent trained on Qwen2.5-7B. Similar to the RAG agent, it uses OpenAI’s `text-embedding-3-large` as the retriever and retrieves the top-8 most relevant chunks. The agent then runs an iterative ReAct loop: generating a search query, retrieving context, reasoning over the results, and deciding whether to search again or answer.
- **MemAgent:** We utilize the official implementation of MemAgent (Yu et al., 2025). MemAgent processes long documents in fixed-size segments (set to 5k tokens by default in the official implementation). It maintains a global memory panel and employs a learnable policy to decide whether to read, write, or overwrite information at each step. We align the prompt settings and base model (Qwen2.5-7B-Instruct) with LycheeMemory to ensure a fair comparison of the memory mechanisms.

A.5 Storage and Computation Trade-off

A critical challenge in long-context processing is the management of GPU memory (VRAM). Even with our efficient compression mechanism, maintaining a compressed KV-cache for extremely long sequences can impose a prohibitive storage overhead.

Storage Analysis. Consider a scenario with a context length of $N = 1.75\text{M}$ tokens using the Qwen2.5-3B model (Hidden size $d = 2304$, Layers $l = 36$). With a compression ratio of $\alpha = 4$, the system generates approximately 437.5k latent memory tokens. Since Qwen2.5-3B utilizes Grouped Query Attention (GQA) with 2 KV heads and 16 Query heads, the storage requirement for the KV-cache (in bfloat16 precision) is:

$$\mathcal{M}_{\text{KV}} \approx 2 \times l \times d_{\text{head}} \times n_{\text{kv}} \times \frac{N}{\alpha} \times 2 \text{ bytes} \approx 18.1 \text{ GB}$$

While this fits within the memory of high-end GPUs like the A100 (80GB), it still consumes a significant portion of VRAM, limiting the space

available for activations and larger batch sizes. Furthermore, our goal is to enable efficient reasoning on consumer-grade hardware and to support scaling to even longer contexts (e.g., 10M tokens), where static storage becomes prohibitive.

Optimization Strategy. To address this, we identify two potential strategies:

- **Offloading:** Temporarily offloading the compressed KV-cache to CPU RAM or NVMe SSDs and swapping them back to GPU only during the retrieval phase.
- **Just-in-Time (JIT) Compression:** Storing only the raw text and using the Compressor to regenerate the latent representations on the fly when needed.

By default, LycheeMemory uses offline pre-compression and stores the compressed KV-cache for inference. For contexts beyond 1M tokens, we optionally enable Just-in-Time (JIT) compression as an engineering strategy to reduce storage overhead. While re-computing embeddings incurs a computational cost, it can be more efficient than the I/O bottleneck of memory swapping. The compression process is a single parallel forward pass, whereas the reasoning process is autoregressive. For a chunk of size 4096, compression requires only 1 forward step. In contrast, generating a reasoning chain often requires hundreds of serial steps. Therefore, the amortized computational overhead of on-the-fly compression is minimal compared to the benefits of reduced memory footprint and improved scalability.

B Training Convergence Analysis

We address the potential concern regarding the stability of jointly optimizing the Compressor and Reasoner, given the sparsity of reward signals in sequence generation tasks. Figure 4 presents the raw, unsmoothed training reward curves over 50 checkpoints, comparing our Joint Optimization strategy against a baseline with a Frozen Compressor.

Observation. End-to-End RL curve (red) shows higher volatility in the early stages (Checkpoints 0-15) compared to the Frozen Compressor (blue). This is expected, as the gradient updates must propagate through the reasoning steps back to the compression module, causing shifts in the memory representation Θ . The Frozen Compressor rapidly converges to a local optimum (≈ 0.67) but fails




Figure 4: Training reward curves (raw data). The blue line (Frozen Compressor) converges quickly but hits a performance plateau. The red line (End-to-End) exhibits higher variance initially due to the exploration of the compression policy but achieving higher rewards.

to improve further, as the reasoner is limited by a static, suboptimal memory bank. In contrast, our method steadily climbs after the initial adaptation phase, reaching a higher reward (≈ 0.7).

Conclusion. The empirical results demonstrate that despite the inherent variance in RL, the joint policy successfully converges. The fluctuating but upward trend confirms that the Compressor is actively learning to retain task-critical features that maximize the Reasoner’s success rate, validating the effectiveness of our end-to-end optimization framework.

C Failure Mode Analysis

To gain deeper insights into the limitations of LycheeMemory, we conducted a manual error analysis on 128 randomly sampled incorrect instances from the HotpotQA and 2WikiMultihopQA validation sets. We categorized the primary causes of failure into three dominant modes: *Compression-Induced Hallucination*, *Unidirectional Dependency Mismatch*, and *Premature Inference Anchoring*.

C.1 Unidirectional Dependency Mismatch (35%)

The most significant source of error (approx. 35%) stems from the inherent limitation of the single-pass, streaming architecture. In multi-hop reasoning, the relevance of an early piece of evidence

often depends on information that appears later in the document.

Mechanism. When the model encounters a critical clue (e.g., at Step 3), it may not be semantically similar to the current query or working memory, causing the Gate to filter it out. Later (e.g., at Step 5), the model discovers the bridge entity that makes the previous clue relevant. However, since the static memory has already been processed and the model cannot backtrack, this information is permanently lost.

Case Study 1: The Late-Binding Problem

Query: *What represents the nationality of the director of the film "The Blue Kite"?*

- Step 3 (Context Chunk): "...Tian Zhuangzhuang was born in Beijing, China, and began his career..."
- Model Action: [Gate: Ignore] → The working memory contains no link to "Tian Zhuangzhuang" yet.
- Step 8 (Context Chunk): "...'The Blue Kite' is a 1993 drama film directed by Tian Zhuangzhuang..."
- Model Action: [Gate: Retrieve] → Update Working Memory: "Director is Tian Zhuangzhuang."
- Reasoning Failure: The model now knows the director, but the information about his nationality (China) was in Step 3, which was discarded. The model answers "Unknown" or hallucinates based on the name.

C.2 Premature Inference Anchoring (21%)

Approximately 21% of errors occur when the model aggressively acts on partial evidence, forming a correct-looking but ultimately wrong conclusion early in the process. This creates a confirmation bias in the Working Memory.

Mechanism. The Reasoner generates an intermediate answer based on a partial match (e.g., a shared name). This incorrect entry in the Working Memory then dominates the attention mechanism, causing the model to either ignore subsequent contradictory evidence or misinterpret it to fit the existing hypothesis.

Case Study 2: Premature Anchoring

Query: Which band's lead singer also released the solo album "Euphoria"?

- Step 2 (Context Chunk): "...Enrique Iglesias released an album titled 'Euphoria' in 2010..."
- Model Action: Update Working Memory: "Candidate: Enrique Iglesias (Solo Artist)." → *Wrong Path. The question asks for a band's lead singer.*
- Step 6 (Context Chunk): "...Def Leppard's lead singer Joe Elliott released a projected titled..." (Irrelevant text follows).
- Step 9 (Context Chunk): "...The band 'Morningwood' features lead singer Chantal Claret..." (Target info appears later).
- Reasoning Failure: The working memory is already anchored on Enrique Iglesias. The model stops actively searching for "bands" or tries to justify why Enrique fits the description, ignoring the correct entity appearing later.

C.3 Compression-Induced Hallucination (17%)

As analyzed in Appendix H.1, about 17% of errors are due to *Feature Collapse* within the Compressor.

Mechanism. High compression ratios can cause distinct entities with similar semantic embeddings (e.g., brothers, movies in the same franchise, dates close in time) to merge in the latent space. The Reasoner retrieves a blurred representation, leading to attribute swapping.

Case Study 3: Attribute Swapping

Query: Who was born earlier, William Johnson or Wilson Johnson?

- Compressed Memory: Encodes "William... born 1856" and "Wilson... born 1860" into adjacent latent vectors.
- Retrieval: The Reasoner retrieves the block containing both.
- Reasoning Failure: Due to vector smoothing, the specific binding of dates to names is lost. The model outputs: "Wilson was born in 1856," effectively swapping the birth years.

C.4 Other Error Types (27%)

The remaining errors include:

- Context Overflow: The accumulation of too many potentially relevant chunks fills the working memory context limit, flushing out early correct evidence.
- Instruction Misalignment: The model correctly retrieves evidence but fails to align the final answer format with user instructions (e.g., answering Yes instead of a specific entity name).

D Computational Complexity

To rigorously evaluate the efficiency of LycheeMemory, we model the theoretical FLOPs required to process a document of length N and generate an answer. We compare three distinct paradigms:

Full-Context. Processes the entire sequence simultaneously.

MemAgent. Processes the sequence in chunks, performing an autoregressive memory update for every chunk.

LycheeMemory (Ours): Compresses the sequence first, then employs a sparse, gated retrieval mechanism.

D.1 FLOPs Formulation

Let N be the total document length, sz be the chunk size, L_Q be the query length, and L_A be the output generation length. The document is segmented into $K = \lceil N/sz \rceil$ chunks. We denote the model's hidden dimension as d and depth as l . The complexity of generating a sequence of length L_{gen} given a prompt of length L_{pmt} is dominated by attention, scaling as $\mathcal{O}(l \cdot d \cdot (L_{pmt} + L_{gen})^2)$.

Full-Context. The computational complexity is dominated by the global self-attention mechanism over the entire sequence.

$$\begin{aligned} C_{\text{Full}} &\approx \mathcal{O}(l \cdot d \cdot (N + L_Q + L_A)^2) \\ &\approx \mathcal{O}(N^2) \quad (\text{when } N \gg L_Q, L_A) \end{aligned}$$

The prefill stage processes $N + L_Q$ tokens, followed by decoding L_A tokens. As N grows to millions, the quadratic term makes this prohibitive.

MemAgent. MemAgent adopts a linear scanning approach but incurs a heavy constant factor due to forced memory updates. It performs generation for *every* chunk. Let the input per step be $L_{\text{in}} = L_Q + \text{Mem}_{\text{size}} + sz$, and output be $L_{\text{out}} = \text{Mem}_{\text{update}}$. The total FLOPs sums over all K steps:

$$\begin{aligned} C_{\text{MemAgent}} &\approx K \times \mathcal{O}(l \cdot d \cdot (L_{\text{in}} + L_{\text{out}})^2) \\ &= \frac{N}{sz} \times C_{\text{gen}} \approx \mathcal{O}(N) \end{aligned}$$

Although asymptotically linear, the constant C_{gen} involves a full generation process (KV-cache read + autoregressive write) for every chunk, leading to a steep increase in computational cost.

LycheeMemory (Ours). Our method decouples processing into efficient compression and sparse reasoning. We utilize a compression ratio of $\alpha = 4$.

- **Phase 1: Compression.** The model processes chunks in parallel to encode KV-cache. Since there is no autoregressive decoding, the cost is proportional to the input tokens.

$$C_{\text{comp}} \approx \mathcal{O}(l \cdot d \cdot N)$$

- **Phase 2: Gate.** The Gate scores all K blocks. Due to compression, the effective sequence length is N/α . The Gate requires only a single forward pass per block.

$$C_{\text{gate}} \approx \mathcal{O}\left(l \cdot d \cdot \frac{N}{\alpha}\right)$$

- **Phase 3: Reasoning.** The Reasoner is activated only for the top- T relevant chunks ($T \ll K$).

$$\begin{aligned} C_{\text{reason}} &\approx T \times \mathcal{O}(l \cdot d \cdot (L_{\text{in}} + L_{\text{out}})^2) \\ &= T \times C_{\text{gen}} \end{aligned}$$

The total cost is $C_{\text{Ours}} = C_{\text{comp}} + C_{\text{gate}} + C_{\text{reason}}$. Comparing the dominant terms:

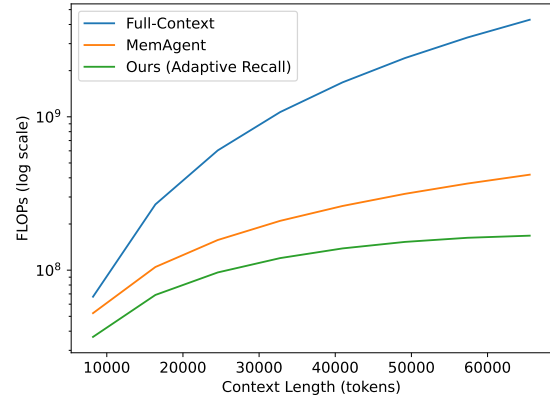


Figure 5: FLOPs Scaling Analysis (Log Scale). We compare the estimated computational cost across context lengths from 8K to 64K tokens using a logarithmic FLOPs axis. For LycheeMemory, the effective recall ratio is assumed to decrease linearly from 100% to 40% as context length increases, reflecting increasingly selective memory access under long contexts. FLOPs are analytically estimated under simplified assumptions and are intended to illustrate relative scaling trends rather than exact runtime measurements.

- **MemAgent:** $\frac{N}{sz} \times C_{\text{gen}}$ (Dense Generation)
- **LycheeMemory:** $\mathcal{O}(N) + T \times C_{\text{gen}}$ (Sparse Generation)

Since $T \ll N/sz$ (e.g., retrieving only 10% of chunks), LycheeMemory significantly reduces the number of expensive generation calls, resulting in a much flatter scaling curve.

D.2 Quantitative Comparison

Figure 5 illustrates the FLOPs scaling behavior under increasing context lengths. The Full-Context baseline exhibits the expected $\mathcal{O}(N^2)$ complexity due to dense self-attention, which appears as a straight line with a steep slope under the logarithmic FLOPs axis, indicating rapidly growing computational cost. In contrast, both MemAgent and LycheeMemory achieve linear scaling with respect to context length. Despite sharing the same asymptotic complexity, LycheeMemory consistently incurs lower FLOPs than MemAgent across all evaluated settings. This improvement is primarily attributed to the **Compress-then-Reason** paradigm: first, the input sequence is compressed by $4\times$, substantially reducing the number of tokens processed by the gating mechanism; second, unlike MemAgent which performs mandatory write operations for every chunk, LycheeMemory employs adaptive gating to activate the computationally heavy Rea-

Model	GovReport				MultiNews			
	R-1	R-2	R-L	Avg.	R-1	R-2	R-L	Avg.
Qwen2.5-7B-Instruct	30.91	11.68	15.20	19.26	46.64	12.01	28.33	28.99
MemAgent-7B	30.28	12.37	15.37	19.34	48.49	14.41	30.91	31.27
LycheeMemory-7B (ours)	30.12	13.08	15.07	19.42	47.43	15.28	30.33	31.01

Table 5: Zero-shot performance on LongBench summarization tasks. Despite being trained for QA, LycheeMemory achieves competitive performance compared to MemAgent, demonstrating strong generalization capabilities.

soner only for a small fraction of chunks, resulting in a significantly smaller constant factor in practice.

E Out-of-Distribution (OOD) Generalization Analysis

While LycheeMemory is primarily optimized for multi-hop QA, we also test whether its compressed memory bank Θ transfers to other long-context formats. We conduct OOD experiments on long-document summarization tasks from LongBench (Bai et al., 2024), which differ from the QA-based training setup. We consider GovReport (summarizing lengthy government reports) and MultiNews (summarizing multiple news documents). We report ROUGE-1 (R-1), ROUGE-2 (R-2), and ROUGE-L (R-L) scores, and compare our method against the base model Qwen2.5-7B-Instruct as well as MemAgent-7B.

Results and Discussion. As shown in Table 5, LycheeMemory exhibits strong zero-shot generalization. These results suggest that the compressed memory bank Θ constructed by LycheeMemory captures transferable semantic features that support both targeted extraction (QA) and global aggregation (summarization).

F Ablation on Working Memory Capacity

In the dynamic recall and reasoning phase, the capacity of the Working Memory (m) is a critical hyperparameter. Given that we segment documents into chunks of size $L_{\text{chunk}} = 4096$, the working memory capacity determines the maximum buffer size available for the active context before synthesis. We investigate the impact of varying the working memory capacity limit ($L_{\text{WM}} \in \{1024, 2048, 3072, 4096\}$) on the RULER-HQA benchmark across increasing context lengths ranging from 7K to 56K tokens.

Results and Analysis. To ensure fairness and isolate the effect of capacity, we excluded the Gate

WM Capacity (L_{WM})	RULER-HQA Context Length				Avg. Score
	7K	14K	28K	56K	
1024	82.03	79.69	78.91	77.34	79.49
2048	82.03	78.91	78.91	77.34	79.30
3072	81.25	78.91	77.34	75.78	78.32
4096	80.47	77.34	76.56	75.00	77.34

Table 6: Ablation results on Working Memory Capacity. The standard capacity of 1024 achieves the best performance, indicating that larger buffers do not necessarily improve reasoning.




Figure 6: Evolution of Working Memory length. The length grows initially as evidence is collected but stabilizes at a peak of ≈ 500 tokens around Turn 23. This demonstrates that LycheeMemory effectively manages its context budget without unbounded growth.

mechanism in this ablation, as it was specifically optimized for 4096-token chunks. As shown in Table 6, we observe that performance does not improve with increased working memory capacity; in fact, $L_{\text{WM}} = 1024$ yields the highest average accuracy of 79.49%. The results suggest that expanding the working memory budget beyond the necessary chunk size tends to introduce excessive noise tokens. This accumulated noise distracts the model’s attention mechanism during the final answer synthesis, leading to a degradation in reasoning precision rather than an improvement.

G Dynamic Evolution of Working Memory

To verify the efficiency of our reasoning mechanism, we tracked the token usage of the Working Memory state m across 32 reasoning steps. Figure 6 presents the average length evolution on 128 samples from the 128K context validation set.

Analysis. As shown in Figure 6, the working memory length exhibits a clear saturation behavior rather than monotonically increasing with reasoning steps. While the memory expands during the early stages to accumulate relevant ev-

Query: Distinguish between the wives of William Henry Johnson and Wilson Harold Johnson.

Context Facts:

1. William’s wife: Elizabeth Ann Smith, born **1860**.
 2. Wilson’s wife: Elizabeth Marie Smith, born **1865**.
-

Model Response ($\alpha = 16$, High Compression):

"William’s wife was Elizabeth Ann... Wilson’s wife was Elizabeth Marie... **Both women were born in 1860**, but their names are distinct."

(Error: *Hallucinated Wilson’s wife’s birth year by merging it with William’s wife’s.*)

Model Response ($\alpha = 4$):

"Elizabeth Ann was born in **1860**, while Elizabeth Marie was born in **1865**. They are cousins."

Table 7: Case study on fine-grained information retrieval. The $16\times$ model suffers from attribute merging (hallucination), while $4\times$ retains precision.

idence, it gradually stabilizes at approximately 503 tokens—well below the predefined capacity limit of 1024. This indicates that the model does not passively append retrieved information, but instead learns to actively regulate its memory state by retaining only task-critical content. Such learned memory management effectively prevents unbounded context growth and ensures stable, efficient reasoning over long interaction horizons.

H Case Analysis

H.1 Hallucination via Feature Collapse

To understand the failure mode of high compression, we conduct a qualitative analysis using a constructed synthetic narrative, *Chronicle of the Four Johnson Brothers*, which is dense with confounding entities (similar names and dates). We compare the generated responses of the $4\times$ (Ours) and $16\times$ models.

Discussion. As shown in Table 7, while the $16\times$ model correctly retrieves high-level entities (names), it fails at *attribute binding*. It incorrectly assigns the birth year “1860” to both wives. This suggests that aggressive compression causes *feature collapse* in the latent space, where distinct numerical tokens (1860 vs. 1865) closer in proximity are averaged into a single representation. In multi-step reasoning, these small hallucinations accumulate, leading to the rapid performance decay observed in the $16\times$ curve (Figure 3). The $2\times$, $4\times$, and $8\times$ models correctly identified the dates in this test, confirming that moderate compression preserves the fidelity required to distinguish fine-grained details.

H.2 Associative Reasoning and Self-Correction

Question: The city that contains the Yunnan Provincial Museum is also known by what nickname?

Chunk k

... The Yunnan Provincial Museum is located in Kunming, the capital city of Yunnan Province in southwest China. The museum features exhibitions on Yunnan’s history, ethnic groups, and archaeology. ...

Working memory k:

- The Yunnan Provincial Museum is located in Kunming City, Yunnan.
- Yangzonghai Lake is located near Kunming, Yunnan, in southern China.

Chunk k+1

... Dian Lake is a large freshwater lake on the Yunnan–Guizhou Plateau near Kunming. Dian Lake is sometimes described as the “Sparkling Pearl Embedded in a Highland”. ...

Working memory k+1:

- The Yunnan Provincial Museum is located in Kunming City, Yunnan.
- Dian Lake is a large lake located on the Yunnan–Guizhou Plateau close to Kunming, Yunnan, its nickname is “Sparkling Pearl Embedded in a Highland”.
- Yangzonghai Lake is located near Kunming, Yunnan, in southern China.

The city that contains the Yunnan Provincial Museum might be “Sparkling Pearl Embedded in a Highland”.

Chunk k+2

... Kunming has a mild climate throughout the year and is widely known as the “Spring City”. The city serves as a major cultural and transportation hub in Yunnan. ...

Working memory k+2:

- The Yunnan Provincial Museum is located in Kunming City, Yunnan.
- Dian Lake is a large lake located on the Yunnan–Guizhou Plateau close to Kunming, Yunnan, its nickname is “Sparkling Pearl Embedded in a Highland”.
- Kunming is also known as the “Spring City” due to its weather.
- Yangzonghai Lake is located near Kunming, Yunnan, in southern China.

The new content provides Kunming’s exact nickname. Therefore, the city that contains the Yunnan Provincial Museum is known as the “Spring City”.

As illustrated in the case study above, the model demonstrates robust *associative reasoning* and *self-correction* capabilities. Initially, the agent identifies “Kunming” as the pivotal entity by associating the Yunnan Provincial Museum with its location (Chunk k). It then forms a tentative hypothesis after processing memory block $k + 1$, incorrectly inferring that the city’s nickname might be “Sparkling Pearl” based on the proximal descriptions of Dian Lake. However, the reasoning remains flexible;

upon retrieving memory block $k + 2$ which explicitly describes Kunming as the "Spring City," the agent successfully detects the conflict. It differentiates the distraction (the lake's nickname) from the city's actual alias and rectifies its working memory, effectively overriding the previous tentative inference with the verified fact.

I Comparison with RAG

RAG and LycheeMemory address different bottlenecks in long-context reasoning. RAG enables fast retrieval with high recall via approximate nearest neighbor search, making it well-suited for large external corpora. Our goal is not to replace RAG, but to provide an alternative long-context processing paradigm based on compressed memory and state-dependent retrieval: RAG typically ranks chunks by query-only similarity and may miss late-hop evidence that becomes relevant only after intermediate entities are discovered, while LycheeMemory conditions the Gate on both the query and the evolving working memory \mathbf{m} (Table 1). The two approaches are complementary: documents retrieved by RAG can be treated as additional context streams, compressed into Θ , and then reasoned over by the same dynamic recall and reasoning workflow.