

LycheeCluster: Efficient Long-Context Inference with Structure-Aware Chunking and Hierarchical KV Indexing

Dongfang Li¹, Zixuan Liu², Gang Lin¹, Baotian Hu¹, and Min Zhang¹

¹Harbin Institute of Technology, Shenzhen

²University of Electronic Science and Technology of China

lidongfang@hit.edu.cn, liuzx@std.uestc.edu.cn

Abstract

The quadratic complexity of the attention mechanism and the substantial memory footprint of the Key-Value (KV) cache present severe computational and memory challenges for Large Language Models (LLMs) processing long contexts. Existing retrieval-based methods often compromise semantic integrity through fixed-size chunking and suffer from inefficient linear scanning. In this paper, we propose LycheeCluster, a novel method for efficient KV cache management. LycheeCluster preserves local semantic coherence via boundary-aware chunking and constructs a recursive hierarchical index rooted in the triangle inequality. This design transforms cache retrieval from a linear scan into a theoretically bounded, logarithmic-time pruning process, while a lazy update strategy supports efficient streaming generation. Experiments demonstrate that LycheeCluster achieves up to a $3.6\times$ end-to-end inference speedup with negligible degradation in model performance, outperforming state-of-the-art KV cache management methods (e.g., Quest, ClusterKV).

1 Introduction

The context window of Large Language Models (LLMs) has expanded dramatically, evolving from 4K to over 2M tokens (Yang et al., 2025a; Comanici et al., 2025; Anthropic, 2025), enabling transformative applications in long-document understanding and complex reasoning. However, utilizing this massive context introduces a critical latency bottleneck. During the autoregressive decoding phase, the attention mechanism requires scanning the entire Key-Value (KV) cache history for every generated token (Tang et al., 2024; Yang et al., 2025b; Tang et al., 2025). As the sequence length grows, the overhead of loading these massive KV tensors from memory consumes substantial bandwidth, causing the decoding speed to deteriorate

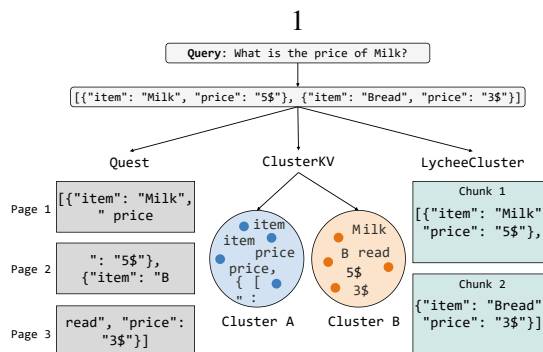


Figure 1: **Impact of retrieval granularity on semantic integrity.** We illustrate the limitations of existing methods using a JSON retrieval query. **(Left) Quest** employs fixed-size pages that arbitrarily sever semantic boundaries. **(Middle) ClusterKV** uses token-level clustering that scatters locally coherent tokens into disjoint clusters based on vector variance. **(Right) LycheeCluster** (Ours) preserves the complete semantic unit via structure-aware chunking, ensuring precise retrieval.

significantly despite the availability of powerful GPUs (Liu et al., 2025b).

To address this memory bandwidth bottleneck without permanently discarding information, *Retrieval-based* methods have emerged as a promising direction (Li et al., 2024a; Liu et al., 2024). Unlike *Eviction-based* methods (e.g., H2O (Zhang et al., 2023), StreamingLLM (Xiao et al., 2024)) that permanently delete tokens, which often leads to irreversible information loss, retrieval-based approaches preserve the full history and dynamically fetch only a small subset of relevant KV pairs for the attention computation. This paradigm aims to accelerate decoding by transforming the linear-complexity attention mechanism into a sparse operation bounded by a fixed token budget.

However, existing retrieval methods struggle with the granularity of their selection strategies, as conceptually illustrated in Figure 1. *Page-based* methods like Quest (Tang et al., 2024) manage KV pairs in fixed-size blocks (e.g., 64 tokens). This rigid segmentation leads to internal fragmentation,

where an entire page is retrieved for a single relevant token, wasting the limited retrieval budget. Conversely, *token-level clustering* methods like ClusterKV (Liu et al., 2025a) aggregate tokens based on global semantic similarity. While this reduces fragmentation, treating tokens as isolated vectors disrupts the local structural integrity of the text. For instance, semantically coupled sequences such as code blocks or reasoning steps are scattered across different clusters, preventing the attention head from retrieving the contiguous context necessary for precise reasoning. Furthermore, relying on global clustering algorithms incurs a high update overhead, often forcing the system to rely on stale indices during streaming generation.

In contrast, we hypothesize that the optimal unit for efficient retrieval is neither the arbitrary fixed-size page nor the isolated token, but the *semantically coherent chunk*. To validate this, we conducted a pilot study (§3) on the StrucText-Eval (Gu et al., 2025). By simply using boundary-aware chunks instead of fixed-size pages, we observed a substantial accuracy improvement (e.g., +15.0%). This finding suggests that preserving the semantic integrity of retrieval units is the key to maximizing the utility of the sparse attention budget.

Inspired by this, we propose LycheeCluster, a novel retrieval-based KV cache management method designed to accelerate long-context decoding through structure-aware indexing. Unlike previous works, LycheeCluster actively segments the context into variable-length chunks based on natural semantic boundaries (e.g., punctuation, line breaks). To enable sub-linear retrieval complexity, we organize these chunks into the **hierarchical indices** (coarse units \rightarrow fine clusters \rightarrow chunks), allowing the model to rapidly prune irrelevant branches using a mathematical upper bound of attention scores. Furthermore, to adapt to the streaming nature of decoding, we further introduce a lazy incremental update strategy. During decoding, new chunks are dynamically grafted onto the nearest existing clusters, maintaining index freshness with negligible runtime overhead.

Extensive experiments on long-context understanding and complex reasoning tasks demonstrate that LycheeCluster effectively breaks the latency bottleneck without compromising the capability. Specifically, our method achieves up to $3.6\times$ decoding speedup compared to full attention, significantly reducing inference latency for long sequences. Crucially, this efficiency does

not come at the cost of precision; LycheeCluster maintains performance comparable to full attention even under aggressive compression budgets. LycheeCluster exhibits exceptional robustness in RULER (Hsieh et al., 2024) and structured data tasks of LONGBENCH V2 (Bai et al., 2025), outperforming token-clustering baselines (e.g., ClusterKV) that often disrupt semantic boundaries. Furthermore, in reasoning-intensive scenarios of MATH500 (Lightman et al., 2023), LycheeCluster successfully supports the dynamic retrieval needs of chain-of-thought generation, preserving logical coherence where heuristic-based baselines often fall.

Our contributions are summarized as follows:

- We identify that structural disruption is a key limitation in existing sparse attention methods and validate structure-aware chunking as a critical solution.
- We propose LycheeCluster, a method featuring hierarchical KV indexing and lazy updates, designed to accelerate inference without compromising semantic integrity.
- We achieve state-of-the-art performance with lower latency among retrieval-based methods, offering a robust solution for deploying long-context and reasoning-intensive LLMs.

2 Related Work

Eviction-based Sparse Attention. To address the memory bottleneck of the KV cache, eviction-based methods reduce memory footprint by permanently discarding tokens based on accumulation metrics or positional heuristics (Zhang et al., 2023; Xiao et al., 2024; Li et al., 2024b). While these methods lower latency, the irreversible deletion of tokens leads to information loss, degrading performance in long-context tasks such as multi-step complex reasoning that require full historical recall (Li et al., 2025; Kim et al., 2025).

Retrieval-based Sparse Attention. Retrieval-based approaches preserve the complete KV history and dynamically fetch only relevant subsets for computation (Liu et al., 2024; Tang et al., 2025; Sun et al., 2025a). For example, Quest (Tang et al., 2024) employs a page-based retrieval strategy, utilizing min-max key statistics of fixed-size blocks to estimate importance. ClusterKV (Liu et al., 2025a) adopts a global perspective, grouping tokens into

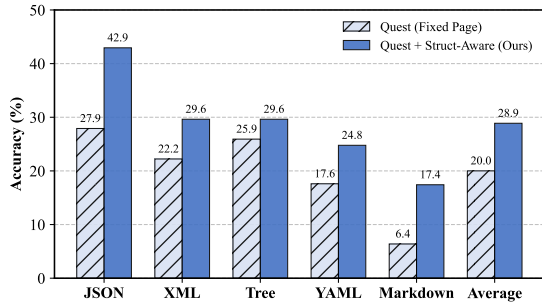


Figure 2: **Pilot Study on StrucText-Eval.** We compare the standard Quest (fixed page) with a modified version using structure-aware chunks while keeping the scoring metric identical. The significant accuracy gain (e.g., **+15.0%** on JSON) confirms that preserving semantic integrity is a prerequisite for effective retrieval.

clusters based on vector similarity to accelerate retrieval. RazorAttention (Tang et al., 2025) further compresses KV cache by selectively caching tokens in different attention heads.

Semantic Granularity and Integrity. A critical limitation in existing sparse attention mechanisms is the disruption of semantic units (Fountas et al., 2025; Zhu et al., 2025). Fixed-size paging arbitrarily severs syntactic boundaries, while token-level clustering fragments locally coherent sequences, such as reasoning steps, into disjoint groups (Chen et al., 2025). Recently, SentenceKV (Zhu et al., 2025) proposed utilizing natural sentences as the atomic unit for caching. However, their rigid punctuation-based segmentation struggles with structured data (e.g., code, JSON) lacking clear sentence delimiters and suffers from length variance. In contrast, we employ a robust structure-aware chunking strategy that respects semantic boundaries while enforcing size constraints, ensuring stable retrieval efficiency across diverse domains.

3 Motivation

Current research on sparse attention primarily focuses on optimizing *importance estimation*, i.e., developing better metrics to decide *which* tokens to keep, such as attention accumulation or min-max statistics. However, a fundamental question remains largely unexplored: **what is the optimal granularity for the atomic unit of scoring?**

Most existing methods adopt rigid segmentation strategies. Quest (Tang et al., 2024) uses fixed-size pages for hardware efficiency, effectively treating the context as a contiguous tape cut at arbitrary intervals. Conversely, ClusterKV (Liu et al., 2025a)

operates on individual tokens grouped by vector similarity, treating the context as a “bag-of-words”. As conceptually illustrated in Figure 1, we hypothesize that both approaches compromise semantic integrity: fixed pages sever context at physical boundaries, while token clustering shatters local coherence.

3.1 A Pilot Study on StrucText-Eval

To validate the hypothesis that granularity matters as much as scoring, we conduct a controlled pilot study using the StrucText-Eval benchmark (Gu et al., 2025). This dataset evaluates LLMs on structured data (e.g., JSON parsing, code completion) where preserving syntactic boundaries is crucial for correctness. We utilize Quest as our experimental baseline due to its representative page-based retrieval mechanism. To isolate the impact of granularity, we kept Quest’s importance estimation algorithm (min-max key approximation) unchanged and modified only the segmentation strategy:

- **Baseline (fixed granularity):** standard implementation using fixed pages of size $S = 16$.
- **Variant (structure-aware granularity):** we replace fixed pages with variable-length chunks defined by natural delimiters (e.g., line breaks, see Table 4). The average chunk size matched baseline to ensure a fair comparison.

Observations. The results are summarized in Figure 2. Notably, merely aligning the retrieval units with semantic boundaries yielded a substantial performance boost. On the StrucText-Eval suite, the structure-aware variant outperformed the original Quest by 10.6% in average accuracy, with gains as high as 15.0% on strictly formatted tasks like JSON extraction.

3.2 Implication

This finding highlights a critical insight: the bottleneck in long-context retrieval is often not the failure to *identify important regions*, but *the fragmentation of those regions*. When a fixed page cuts a logical unit (e.g., a function definition) in half, or when clustering scatters a KV cache pair, the attention head fails to attend to the complete information even if the retrieval score is high. This **semantic misalignment** problem renders the retrieved KV cache pairs partially useless.

Therefore, simply patching existing methods is insufficient. We need a retrieval framework designed *natively* around the **structure-aware chunk**

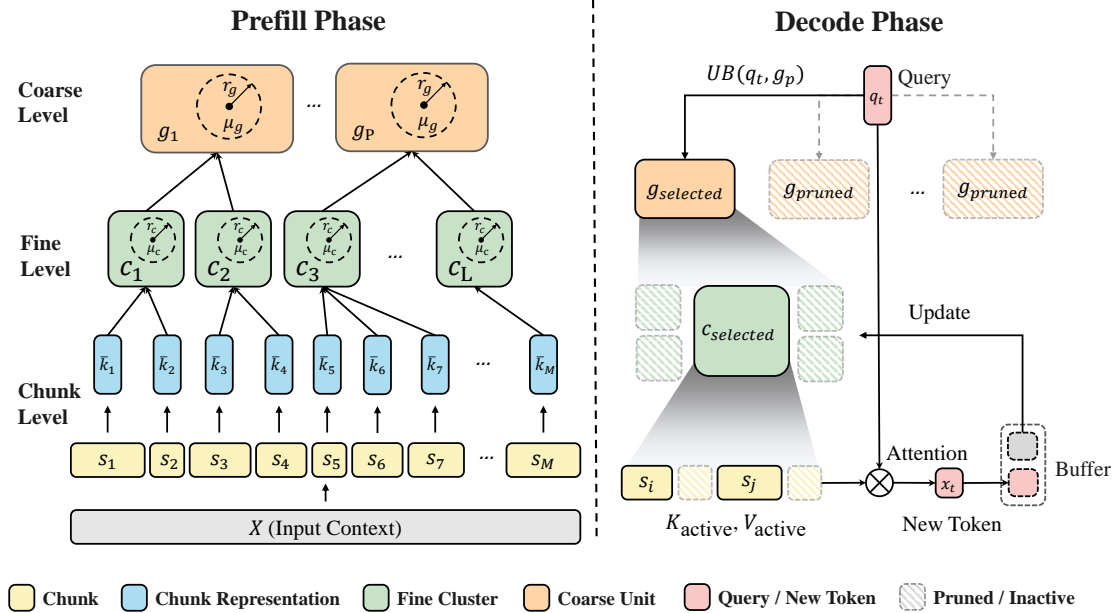


Figure 3: The overall pipeline of LycheeCluster. The left panel illustrates the bottom-up index construction during the prefill phase, where variable-length chunks are hierarchically clustered. The right panel demonstrates the top-down retrieval and incremental update during the decoding phase.

as the atomic unit. This motivation drives the design of **LycheeCluster**, which we formally introduce in the next section. By enforcing semantic atomicity at the indexing level, we ensure that every retrieved unit provides a complete, actionable context for the attention mechanism.

4 Method

In this section, we describe LycheeCluster in detail, a retrieval-based approach designed to accelerate long-context LLM inference.

4.1 Formulation

Given a Transformer model \mathcal{M} and an input sequence $X = \{x_1, \dots, x_{t-1}\}$, the model generates the next token at step t via the attention mechanism. Let $q_t \in \mathbb{R}^d$ be the query vector; the output is a weighted sum of historical KV pairs:

$$\text{Attn}(q_t, K_{<t}, V_{<t}) = \sum_{i=1}^{t-1} \frac{\exp(q_t^\top k_i / \sqrt{d})}{Z} v_i \quad (1)$$

where Z is the normalization factor. Naively computing $q_t^\top k_i$ for all historical tokens leads to linear complexity $\mathcal{O}(t)$. Our objective is to approximate this process by retrieving only the most relevant subset of KV pairs.

As a step towards this, we construct a recursive hierarchical abstraction. Rather than index individual tokens, we organize the KV cache into a

tree structure \mathcal{T} . First, we partition the token sequence into a set of semantically coherent chunks $\mathcal{S} = \{s_1, \dots, s_M\}$. For each chunk s_j , we define its representative key \bar{k}_j by performing mean pooling over its constituent token keys and projecting the result onto the unit sphere ($\|\bar{k}_j\|_2 = 1$). The premise is that the strong local coherence of semantic information in sequences, the relevance of a chunk can be effectively approximated by the relevance of its aggregated representation (Lu et al., 2025). Second, to facilitate efficient search, we build a multi-level index on top of \mathcal{S} . We define a parent node u (e.g., a cluster centroid) that covers a set of child nodes \mathcal{V}_u (e.g., chunks or sub-clusters). For each node u , we maintain its centroid μ_u and a covering radius $r_u = \max_{v \in \mathcal{V}_u} \|v - \mu_u\|_2$. Third, leveraging the triangle inequality and the Cauchy-Schwarz inequality, we derive a strict upper bound for the relevance score. For any query q_t and any child node v within the parent node u , the dot product similarity satisfies:

$$\begin{aligned} q_t^\top v &= q_t^\top (\mu_u + (v - \mu_u)) \\ &\leq \underbrace{q_t^\top \mu_u}_{\text{Score Upper Bound (UB)}} + \|q_t\|_2 \cdot r_u. \end{aligned} \quad (2)$$

Hence, Eqn. 2 serves as the theoretical foundation for our pruning strategy: the similarity between the query and a centroid, plus a slack term determined by the radius, strictly bounds the maximum possible similarity for any element within

Algorithm 1 Inference Process with LycheeCluster

Input: Prompt X , Model \mathcal{M} , Retrieval budgets k_g, k_c
Output: Generated sequence Y

- 1: // Phase 1: Index Construction (Prefill)
- 2: $\mathcal{S} \leftarrow \text{StructureAwareChunking}(X)$
- 3: $\mathcal{I} \leftarrow \text{HierarchicalIndexing}(\mathcal{S}) \triangleright$ Build Coarse \mathcal{G} & Fine \mathcal{C}
- 4: $B \leftarrow \emptyset$ \triangleright Initialize token buffer
- 5: // Phase 2: Decoding & Retrieval
- 6: **for** $t = 1, \dots, T$ **do**
- 7: $q_t \leftarrow \mathcal{M}.\text{GetQuery}(x_t)$
- 8: // Step 1: Coarse-Level Pruning (Top- k_g)
- 9: $\text{Scores}_g \leftarrow \{q_t^\top \mu_g + \|q_t\| \cdot r_g \mid g \in \mathcal{G}\} \triangleright$ Eqn. 2
- 10: $\mathcal{G}_{top} \leftarrow \text{TopK}(\text{Scores}_g, k_g)$
- 11: // Step 2: Fine-Level Pruning (Top- k_c)
- 12: $\mathcal{C}_{cand} \leftarrow \bigcup_{g \in \mathcal{G}_{top}} g.\text{children}$
- 13: $\text{Scores}_c \leftarrow \{q_t^\top \mu_c + \|q_t\| \cdot r_c \mid c \in \mathcal{C}_{cand}\}$
- 14: $\mathcal{C}_{top} \leftarrow \text{TopK}(\text{Scores}_c, k_c)$
- 15: // Step 3: Exact Attention
- 16: $KV_{active} \leftarrow \bigcup_{c \in \mathcal{C}_{top}} c.\text{chunks.KV}$
- 17: $x_{t+1} \leftarrow \text{Attention}(q_t, KV_{active})$
- 18: // Step 4: Incremental Update
- 19: $B.\text{append}(\mathcal{M}.\text{GetKV}(x_t))$
- 20: **if** $|B| \geq \text{ChunkSize}$ **then**
- 21: $s_{new} \leftarrow \text{Pack}(B); B \leftarrow \emptyset$
- 22: $\text{LazyUpdate}(\mathcal{I}, s_{new}) \triangleright$ Assign & expand radius
- 23: **end if**
- 24: **end for**

that cluster. This allows us to safely discard entire branches of the index tree if their upper bound is sufficiently low, without inspecting individual chunks or tokens.

4.2 Overall Framework

As illustrated in Figure 3, the LycheeCluster framework operates in two distinct phases: **(1) index construction (prefill phase)**. When receiving long-context inputs, we do not store a flat KV cache. Instead, we employ a structure-aware algorithm to segment the context, organizing it into a three-level index: coarse units \rightarrow fine clusters \rightarrow chunks. **(2) retrieval & update (decoding phase)**. During token generation, the query q_t traverses the index tree top-down. By calculating the upper bound by Eqn. 2, the system rapidly identifies the most relevant KV chunks. Concurrently, newly generated tokens are managed via a lightweight incremental mechanism to maintain index freshness without incurring heavy re-clustering overheads. We provide the details of LycheeCluster in Algorithm 1.

4.3 Prefill Phase

In the prefill phase, we process the full prompt context to build the retrieval-ready structure. This involves two key components: structure-aware chunking and hierarchical KV indexing.

Structure-Aware Chunking. Standard fixed size window approaches often disrupt semantic dependencies or syntactic structures, degrading representation quality. To mitigate this, we propose a boundary-aware segmentation algorithm. The algorithm accumulates tokens greedily and, upon reaching a minimum length threshold, looks ahead for high-priority natural delimiters (e.g., paragraph breaks, sentence endings, punctuation). If the maximum threshold is reached without a natural break, a forced split is applied. This process maps the context into a sequence of variable-length, semantically self-contained chunks $S = \{s_1, s_2, \dots, s_M\}$.

Hierarchical KV Indexing. To enable logarithmic time retrieval, we organize the chunks into a pyramid structure. The process begins at the chunk level, where we compute the representation \bar{k}_i for each chunk s_i by performing mean pooling on its internal token keys followed by L_2 normalization. Building upon these representations, we further proceed to the fine cluster level by applying spherical k -means clustering (Hornik et al., 2012) to partition $\{\bar{k}_i\}$ into L fine-grained clusters. Each cluster c_j is characterized by its centroid μ_{c_j} and a covering radius r_{c_j} , defined as the maximum Euclidean distance from the centroid to any member chunk. Finally, to address scenarios with extremely long contexts where L remains prohibitively large, we construct a coarse unit level. We employ the same logic to aggregate the L cluster centroids into P coarse units ($P \ll L$). This three-tier architecture (coarse \rightarrow fine \rightarrow chunk) drastically reduces the search space for the subsequent decoding phase.

4.4 Decoding Phase

During decoding, the system must locate high-relevance caches with sub-linear complexity while accommodating the growing context.

Top-Down Pruning. For a query q_t at each step, we perform a hierarchical search. We first evaluate the upper bound (Eqn. 2) against the top-level coarse units. Specifically, we compute $\text{UB}(q_t, g_p) = q_t^\top \mu_{g_p} + \|q_t\| \cdot r_{g_p}$ and retain only the top- k_g units. Within these selected units, we apply the same logic to rank and select the top- k_c fine clusters. Finally, all KV chunks belonging to the selected fine clusters are loaded as the active set (K_{active}, V_{active}) for exact attention computation. This *prune-and-refine* strategy ensures that the computational cost remains nearly constant regardless of total context length.

Incremental Update. As new tokens are generated, their corresponding KV cache pairs must be indexed. To avoid the prohibitive cost of global re-clustering, LycheeCluster adopts a lazy update strategy. During decoding, new tokens are temporarily stored in a buffer. Once enough tokens have accumulated in the buffer to form a full *dynamic chunk*, the chunk is assigned to the nearest existing fine cluster and coarse unit based on centroid proximity. The centroids of the affected nodes are updated via moving averages, while the radii undergo monotonic expansion to encompass the new chunk. This design allows the model to handle infinite streaming generation while maintaining the integrity of the retrieval index.

5 Experiments

5.1 Experimental Settings

Benchmarks and Models. To evaluate the performance over different scenarios, we conduct experiments across models of varying architectures and sizes. For long-context understanding, we benchmark the Llama-3.1-8B-Instruct (Team, 2024) model on the LONGBENCH V2 (Bai et al., 2025) dataset. It contains complex real world long-context questions across diverse tasks such as long in-context learning, long structured data analysis, and code repository understanding. With context lengths ranging from 8K to 2M, LONGBENCH V2 emphasizes deep reasoning and noise resistance, distinguishing it from LongBench (Bai et al., 2024) which focuses on simple retrieval. This allows for rigorous testing of semantic understanding over extremely long sequences. For complex reasoning scenarios, we employ the Deepseek-R1-Distill-Llama-8B and Deepseek-R1-Distill-Qwen-14B (DeepSeek-AI, 2025) models, evaluating them on the MATH500 (Hendrycks et al., 2021a) dataset. It consists of challenging mathematics competition problems and is widely used to assess long chain-of-thought reasoning performance. Due to space limitations, results of RULER benchmark (Hsieh et al., 2024) on Appendix H.

Baseline. We compare LycheeCluster against full attention and several state-of-the-art KV management and sparse attention methods including Quest (Tang et al., 2024), ClusterKV (Liu et al., 2025a), ArkVale (Chen et al., 2024), RaaS (Hu et al., 2025), ShadowKV (Sun et al., 2025a), and RazorAttention (Tang et al., 2025). Since ClusterKV is equivalent to full attention when context

length does not exceed the budget, and the context lengths in MATH500 are relatively short, we exclude it from complex reasoning experiments. The implementation details are in Appendix A.

5.2 Performance Evaluation

Performance on Long-Context Understanding. As shown in Table 1, the results indicate that on the LongBench V2 benchmark, LycheeCluster achieves an overall accuracy of 30.8% with a 1024 token budget. The margin compared with other methods validates that preserving semantic integrity via structure-aware chunking is superior to rigid paging or token-level clustering, which often fragment context. Furthermore, the slight improvement over full attention suggests that our hierarchical pruning effectively acts as a noise filter, enabling the model to focus on critical information within noisy long contexts.

Performance on Complex Reasoning. We evaluate the reasoning capability of LycheeCluster on the MATH500, results shown in Table 2. Complex mathematical reasoning typically involves long chains of thought, where recalling early premises is crucial. LycheeCluster keeps performance loss within 2% on Deepseek-R1-Distill-Llama-8B, and surpasses the full attention on Deepseek-R1-Distill-Qwen-14B. This demonstrates that our method effectively adapts to the dynamic KV cache distribution; even after generating thousands of tokens, it can rapidly retrieve key information from the problem statement and reasoning process, effectively overcoming the bottlenecks of traditional sparse attention in complex reasoning.

5.3 Efficiency Evaluation

End-to-End Decoding Speedup. We use Time Per Output Token (TPOT) as the primary metric to evaluate the decoding latency of LycheeCluster, comparing it with ClusterKV and full attention based on FlashAttention-2 (Dao, 2024)) under varying context lengths and token budgets. As illustrated in Figure 4, the decoding latency of full attention grows linearly with context length. In comparison, LycheeCluster consistently maintains low latency. At a 32K context length, we achieve a 2.6 \times speedup relative to full attention; when the context extends to 64K, this speedup reaches 3.6 \times . This result confirms the high efficiency of LycheeCluster when processing long contexts.

Method	Overall	Short	Medium	Long
Full Attention	30.02	35.56	27.91	25.00
RazorAttention (Tang et al., 2025)	27.44	33.89	25.12	21.30
RaaS (Hu et al., 2025)	28.23	33.89	26.51	22.02
ArkVale (Chen et al., 2024)	28.63	33.89	26.98	23.15
ShadowKV (Sun et al., 2025a)	25.45	32.78	22.79	18.52
Quest (Tang et al., 2024)	20.68	22.78	19.07	20.37
ClusterKV (Liu et al., 2025a)	26.64	30.56	24.65	24.07
LycheeCluster	30.82	37.22	28.84	24.07

Table 1: Performance comparison on LongBench V2 using Llama-3.1-8B-Instruct. **Bold** denotes the best sparse method; **Blue** highlights cases where LycheeCluster outperforms Full Attention.

Method	DeepSeek-R1-Distill	DeepSeek-R1-Distill
	Llama-8B	Qwen-14B
Full Attention	78.40	74.00
RazorAttention	72.00	70.00
RaaS	74.00	68.00
ArkVale	72.00	72.00
ShadowKV	76.00	76.00
Quest	72.00	76.00
LycheeCluster	77.00	74.80

Table 2: Accuracy on MATH500. LycheeCluster demonstrates consistent improvements over sparse baselines across different model architectures on DeepSeek-R1-Distill-Llama-8B.

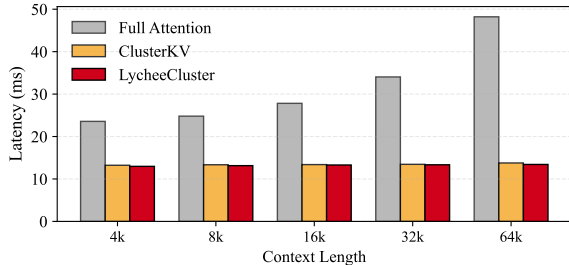


Figure 4: End-to-end decoding latency (TPOT) comparison on a single H20 GPU across varying context lengths. While full attention exhibits linear latency growth, our method maintains consistently low latency.

Kernel-Level Analysis. To gain deeper insights into the sources of efficiency improvements and justify the overhead of clustering and retrieval, we decompose the runtime into four components: (1) index construction (prefill only), (2) hierarchical retrieval, (3) index update and (4) sparse attention.

Figure 5(a) shows the time breakdown during the prefill phase. Although constructing the hierarchical index introduces additional computation, measurements indicate that this process accounts for only 10% - 15% of the total prefill time. Considering that prefill is a one-time operation and subsequent decoding typically involves generating

thousands of tokens, its amortized cost is almost negligible. In contrast, Figure 5(b) breaks down the latency of a single decoding step at a 72K context length, with the total time divided into hierarchical retrieval, index update and sparse attention. Because of the logarithmic complexity of hierarchical retrieval, the retrieval operation occupies only a minimal fraction of the single-step decoding time; this overhead is far outweighed by the benefits of reduced attention computation. Furthermore, the cost of index maintenance is negligible. Our lazy update strategy consumes less than 1% of the decoding time. These results indicate that LycheeCluster achieves high-precision retrieval without disrupting generation fluency.

5.4 Ablation Study

Impact of Structure-Aware Chunking. To assess the impact of our structure-aware chunking design, we replace the structure-aware chunking with fixed-size chunking (chunk size = 16) while keeping the rest of the method unchanged. As illustrated in Figure 6, this results in a 3.03% performance drop for LycheeCluster on the Long Structured Data Understanding task of LongBench V2, along with a degradation in Code Repository Understanding. This confirms that for semantic retrieval to function effectively, the fundamental unit of retrieval (chunk) must be semantically complete.

Pooling Strategy for Chunk Representation. We investigate the impact of the aggregation function used to generate the representative key \bar{k}_j for each chunk. We compare our mean pooling strategy against max pooling. The results are illustrated in Table 3, mean pooling consistently outperforms max pooling across all subsets. We attribute this to the geometric nature of the self-attention mechanism. Since attention scores are derived from dot products (measuring alignment between Query and




Figure 5: Kernel-level latency breakdown. **(a) Prefill Phase:** Latency comparison across varying context lengths. The colored top sections represent the index construction overhead. While LycheeCluster incurs a slightly higher construction cost (10–15%) than ClusterKV, it remains a minor fraction of the total prefill time. **(b) Decoding Phase:** Breakdown of total latency for generating 1,024 tokens at 72k context. The combined overhead of retrieval and lazy updates in LycheeCluster is minimal compared to the massive computation reduction.

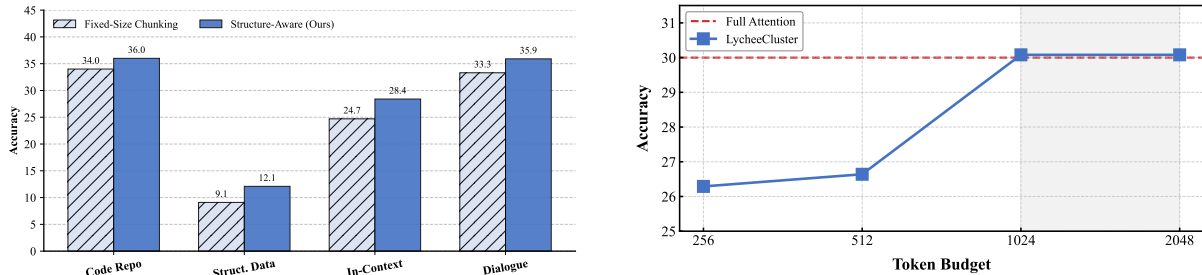


Figure 6: Ablation study on the impact of segmentation strategies. We compare the accuracy of LycheeCluster using our structure-aware chunking against a fixed-size chunking baseline across four task categories. The results demonstrate that preserving semantic boundaries consistently improves performance.

Strategy	Overall	Short	Medium	Long	Recall Rate
Max	28.83	35.56	26.98	21.3	33.60%
Mean	30.82	37.22	28.84	24.07	40.37%

Table 3: Ablation study on representative key pooling strategies. We compare mean pooling (Ours) against max pooling on the LongBench V2 (Llama-3.1-8B-Instruct). **Recall Rate** denotes the retrieval recall rate. It is defined as the percentage of the top- k tokens with the highest ground-truth attention scores (computed by full attention) that are successfully retrieved by the strategy within the same token budget.

Key vectors), preserving the directionality of the feature space is crucial. Max pooling constructs a synthetic vector composed of extreme values from each dimension, which often distorts the original semantic direction and is sensitive to outliers. In contrast, mean pooling followed by L_2 normalization effectively computes the geometric centroid of the token embeddings of the chunk. This provides a faithful representation of the average semantic position of the chunk, mathematically aligning with

Figure 7: Impact of token budget on LongBench V2 performance. The accuracy improves as the budget increases from 256 to 1024.

our spherical k -means objective to maximize the retrieval recall of relevant semantic information.

Impact of Token Budget. We analyze the impact of the token budget on performance. As shown in Figure 7, overall accuracy in LongBench V2 improves significantly as the budget increases from 256 to 1024, but then saturates, showing limited gains with larger budgets. This demonstrates that LycheeCluster can effectively utilize a small budget to precisely capture the critical tokens within extremely long contexts that are essential for generation, thereby maintaining model performance.

6 Conclusion

We proposed LycheeCluster that harmonizes memory efficiency with semantic integrity. By employing structure-aware chunking and hierarchical index, we achieve bounded, sub-linear retrieval without disrupting local context. Our approach outperforms state-of-the-art methods, delivering comparable accuracy to full attention with lower latency. It offers a scalable solution for long-context LLMs inference in resource-constrained scenarios.

Limitations

First, one limitation of our method is that we focus on speeding up generation, and the proposed method does not accelerate attention during the prefill phase. However, prefill overhead is a one-time cost. In long-context scenarios (e.g., RAG, agentic workflows), the decoding phase (generation) typically dominates the total latency. The 3.6x speedup in decoding quickly offsets the prefill cost. One promising direction for future research would be to extend our method to support accelerating prefill. Second, integrating our proposed method into established LLM serving frameworks (e.g., vLLM (Kwon et al., 2023), SGLang (Zheng et al., 2024)) necessitates additional engineering effort. Furthermore, more adaptive strategies can further enhance the generalization ability of this method. As part of ongoing and future work, we are actively investigating the relationship between input features (e.g., sentence length distribution and semantic diversity) and the optimal semantic-aware chunking factor. Preliminary observations suggest that dynamically adjusting this factor based on these input statistics can further improve efficiency and accuracy. Finally, while our technique incurs a memory overhead for storing centroids although this cost is negligible when compared to the total size of the KV cache.

Ethical Considerations

Our work primarily contributes to Green AI by reducing the inference energy consumption of LLMs through sub-linear retrieval complexity. This efficiency enables the deployment of long-context models on resource-constrained hardware, democratizing access to advanced AI. However, as an approximate attention mechanism, there is a theoretical risk of information loss in safety-critical domains (e.g., medical or legal analysis), necessitating careful verification in high-stakes applications. Finally, while efficient inference lowers barriers for beneficial use, it could also facilitate large-scale misuse by malicious actors; we advocate for responsible deployment to mitigate these risks.

Acknowledgments

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). Dongfang Li is the corresponding author.

References

- Anthropic. 2025. Claude 3.7 sonnet system card.
- Yushi Bai, Xin Lv, Jiajie Zhang, Hongchang Lyu, Jiankai Tang, Zhidian Huang, Zhengxiao Du, Xiao Liu, Aohan Zeng, Lei Hou, Yuxiao Dong, Jie Tang, and Juanzi Li. 2024. [Longbench: A bilingual, multi-task benchmark for long context understanding](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, ACL 2024, Bangkok, Thailand, August 11-16, 2024, pages 3119–3137. Association for Computational Linguistics.
- Yushi Bai, Shangqing Tu, Jiajie Zhang, Hao Peng, Xiaozhi Wang, Xin Lv, Shulin Cao, Jiazheng Xu, Lei Hou, Yuxiao Dong, and 1 others. 2025. Longbench v2: Towards deeper understanding and reasoning on realistic long-context multitasks. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3639–3664.
- Jinhan Chen, Jianchun Liu, Hongli Xu, Xianjun Gao, and Shilong Wang. 2025. Sablock: Semantic-aware kv cache eviction with adaptive compression block size. *arXiv preprint arXiv:2510.22556*.
- Renze Chen, Zhuofeng Wang, Beiquan Cao, Tong Wu, Size Zheng, Xiuhong Li, Xuechao Wei, Shengen Yan, Meng Li, and Yun Liang. 2024. [Arkvalde: Efficient generative LLM inference with recallable key-value eviction](#). In *Advances in Neural Information Processing Systems 38: Annual Conference on Neural Information Processing Systems 2024, NeurIPS 2024, Vancouver, BC, Canada, December 10 - 15, 2024*.
- 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*.
- Tri Dao. 2024. [Flashattention-2: Faster attention with better parallelism and work partitioning](#). In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*.
- DeepSeek-AI. 2025. [Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning](#). *CoRR*, abs/2501.12948.
- Zafeirios Fountas, Martin Benfeghoul, Adnan Oomeree, Fenia Christopoulou, Gerasimos Lampouras,

- Haitham Bou Ammar, and Jun Wang. 2025. Human-inspired episodic memory for infinite context llms. In *The Thirteenth International Conference on Learning Representations*.
- Zhouhong Gu, Haoning Ye, Xingzhou Chen, Zeyang Zhou, Hongwei Feng, and Yanghua Xiao. 2025. Structext-eval: Evaluating large language model’s reasoning ability in structure-rich text. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 223–244.
- Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. 2021a. **Measuring mathematical problem solving with the MATH dataset**. In *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks 1, NeurIPS Datasets and Benchmarks 2021, December 2021, virtual*.
- Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. 2021b. Measuring mathematical problem solving with the math dataset. *NeurIPS*.
- Kurt Hornik, Ingo Feinerer, Martin Kober, and Christian Buchta. 2012. Spherical k-means clustering. *Journal of statistical software*, 50:1–22.
- Cheng-Ping Hsieh, Simeng Sun, Samuel Kriman, Shantanu Acharya, Dima Rekeshe, 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*.
- Junhao Hu, Wenrui Huang, Weidong Wang, Zhenwen Li, Tiancheng Hu, Zhixia Liu, Xusheng Chen, Tao Xie, and Yizhou Shan. 2025. **Raas: Reasoning-aware attention sparsity for efficient LLM reasoning**. In *Findings of the Association for Computational Linguistics, ACL 2025, Vienna, Austria, July 27 - August 1, 2025*, pages 2577–2590.
- Jang-Hyun Kim, Jinuk Kim, Sangwoo Kwon, Jae W. Lee, Sangdoo Yun, and Hyun Oh Song. 2025. **Kvzip: Query-agnostic KV cache compression with context reconstruction**. *CoRR*, abs/2505.23416.
- Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph E. Gonzalez, Ion Stoica, and Zhanghao Wu. 2023. Efficient memory management for large language model serving with pagedattention. In *Proceedings of the 29th Symposium on Operating Systems Principles*, pages 611–626.
- Haoyang Li, Yiming Li, Anxin Tian, Tianhao Tang, Zhanchao Xu, Xuejia Chen, Nicole Hu, Wei Dong, Qing Li, and Lei Chen. 2024a. A survey on large language model acceleration based on kv cache management. *arXiv preprint arXiv:2412.19442*.
- Yucheng Li, Huiqiang Jiang, Qianhui Wu, Xufang Luo, Surin Ahn, Chengruidong Zhang, Amir H Abdi, Dongsheng Li, Jianfeng Gao, Yuqing Yang, and 1 others. 2025. Scbench: A kv cache-centric analysis of long-context methods. In *The Thirteenth International Conference on Learning Representations*.
- Yuhong Li, Yingbing Huang, Bowen Yang, Bharat Venkitesh, Acyr Locatelli, Hanchen Ye, Tianle Cai, Patrick Lewis, and Deming Chen. 2024b. **Snapkv: LLM knows what you are looking for before generation**. In *Advances in Neural Information Processing Systems 38: Annual Conference on Neural Information Processing Systems 2024, NeurIPS 2024, Vancouver, BC, Canada, December 10 - 15, 2024*.
- Hunter Lightman, Vineet Kosaraju, Yuri Burda, Harrison Edwards, Bowen Baker, Teddy Lee, Jan Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. 2023. Let’s verify step by step. In *The Twelfth International Conference on Learning Representations*.
- Di Liu, Meng Chen, Baotong Lu, Huiqiang Jiang, Zhenhua Han, Qianxi Zhang, Qi Chen, Chengruidong Zhang, Bailu Ding, Kai Zhang, and 1 others. 2024. Retrievalattention: Accelerating long-context llm inference via vector retrieval. *arXiv preprint arXiv:2409.10516*.
- Guangda Liu, Chengwei Li, Jieru Zhao, Chenqi Zhang, and Minyi Guo. 2025a. **Clusterkv: Manipulating LLM KV cache in semantic space for recallable compression**. In *62nd ACM/IEEE Design Automation Conference, DAC 2025, San Francisco, CA, USA, June 22-25, 2025*, pages 1–7. IEEE.
- Jiaheng Liu, Dawei Zhu, Zhiqi Bai, Yancheng He, Huanxuan Liao, Haoran Que, Zekun Wang, Chenchen Zhang, Ge Zhang, Jiebin Zhang, and 1 others. 2025b. A comprehensive survey on long context language modeling. *arXiv preprint arXiv:2503.17407*.
- Enzhe Lu, Zhejun Jiang, Jingyuan Liu, Yulun Du, Tao Jiang, Chao Hong, Shaowei Liu, Weiran He, Enming Yuan, Yuzhi Wang, Zhiqi Huang, Huan Yuan, Suting Xu, Xinran Xu, Guokun Lai, Yanru Chen, Huabin Zheng, Junjie Yan, Jianlin Su, and 6 others. 2025. **MoBA: Mixture of block attention for long-context LLMs**. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*.
- Hanshi Sun, Li-Wen Chang, Wenlei Bao, Size Zheng, Ningxin Zheng, Xin Liu, Harry Dong, Yuejie Chi, and Beidi Chen. 2025a. **ShadowKV: KV cache in shadows for high-throughput long-context LLM inference**. In *Forty-second International Conference on Machine Learning*.
- Yutao Sun, Tianzhu Ye, Li Dong, Yuqing Xia, Jian Chen, Yizhao Gao, Shijie Cao, Jianyong Wang, and Furu Wei. 2025b. Rectified sparse attention. *arXiv preprint arXiv:2506.04108*.
- Hanlin Tang, Yang Lin, Jing Lin, Qingsen Han, Danning Ke, Shikuan Hong, Yiwu Yao, and Gongyi Wang. 2025. Razorattention: Efficient kv cache compression through retrieval heads. In *The Thirteenth International Conference on Learning Representations*.

- Jiaming Tang, Yilong Zhao, Kan Zhu, Guangxuan Xiao, Baris Kasikci, and Song Han. 2024. [QUEST: query-aware sparsity for efficient long-context LLM inference](#). In *Forty-first International Conference on Machine Learning, ICML 2024, Vienna, Austria, July 21-27, 2024*. OpenReview.net.
- Llama Team. 2024. [The llama 3 herd of models](#). *CoRR*, abs/2407.21783.
- Guangxuan Xiao, Yuandong Tian, Beidi Chen, Song Han, and Mike Lewis. 2024. [Efficient streaming language models with attention sinks](#). In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net.
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, and 1 others. 2025a. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*.
- Lijie Yang, Zhihao Zhang, Zhuofu Chen, Zikun Li, and Zhihao Jia. 2025b. [Tidaldecode: Fast and accurate LLM decoding with position persistent sparse attention](#). In *The Thirteenth International Conference on Learning Representations, ICLR 2025, Singapore, April 24-28, 2025*. OpenReview.net.
- Zhenyu Zhang, Ying Sheng, Tianyi Zhou, Tianlong Chen, Lianmin Zheng, Ruisi Cai, Zhao Song, Yuandong Tian, Christopher Ré, Clark W. Barrett, Zhangyang Wang, and Beidi Chen. 2023. [H2O: heavy-hitter oracle for efficient generative inference of large language models](#). In *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023*.
- Lianmin Zheng, Liangsheng Yin, Zhiqiang Xie, Chuyue Livia Sun, Jeff Huang, Cody Hao Yu, Shiyi Cao, Christos Kozyrakis, Ion Stoica, Joseph E Gonzalez, and 1 others. 2024. Sglang: Efficient execution of structured language model programs. *Advances in neural information processing systems*, 37:62557–62583.
- Yuxuan Zhu, Ali Falahati, David H. Yang, and Mohammad Mohammadi Amiri. 2025. [SentenceKV: Efficient LLM inference via sentence-level semantic KV caching](#). In *Second Conference on Language Modeling*.

A Implementation Details

All experiments are conducted on NVIDIA H20 GPUs. We implement a series of high-performance kernels based on CUDA C++ to maximize hardware utilization and memory access efficiency. Building upon ClusterKV (Liu et al., 2025a), we specifically redesigned the logic for index construction and token selection (retrieval). We further design and write custom kernels for variable-length chunk parallel pooling, label assignment, centroid updates, and radius calculations. Regarding the hyperparameters of LycheeCluster, we set the default token budget to 1024. For the chunking strategy, we set the minimum and maximum chunk length thresholds to 8 and 16 tokens, respectively, with a buffer size of 128 tokens. This range is derived from the average information density of natural language sentences and code statements, generally covering a complete semantic unit across most LLMs tokenizers. Unlike rigid fixed-size chunking methods that mechanically cut data and potentially split critical entities (e.g., variable names), our method ensures that atomic components remain intact. Moreover, the algorithm actively searches for the best syntactic breakpoints (e.g., commas or operators) within the effective window to ensure syntactically safe segmentation (Appendix §B). For the clustering mechanism, we employ spherical k -means using the inner product as the distance metric, with the number of iterations fixed at 10. We observe that the initialization and the number of convergence iterations have a negligible impact on the final performance. Specifically, setting the iteration count to 10-20 yields similar results. To determine the initial number of cluster centers, we set the average fine cluster size to 2 chunks, while limiting the maximum number of coarse units to 64. To ensure stability and align with settings in previous work, we retain full key-value pairs for the first 2 layers and set the attention sink size to 16. All baselines were reproduced based on their official open-source code and recommended parameters.

B Details of Natural Delimiters

We define a hierarchical set of separators to guide the recursive splitting process, as detailed in Table 4. The separators are categorized into four priority levels based on their semantic strength. Level 1 identifies major structural boundaries, such as paragraph breaks and syntax-specific delimiters found in Markdown or code. Level 2 preserves sentence

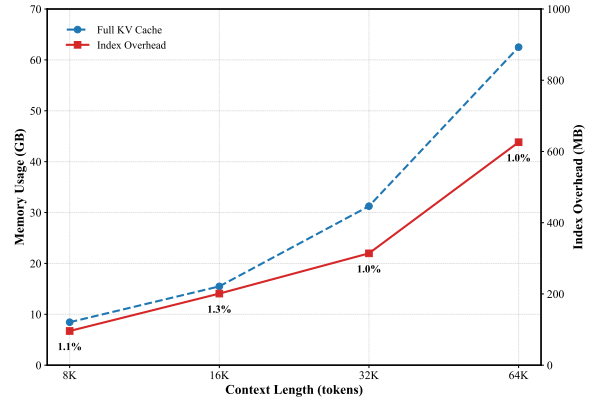


Figure 8: Memory usage comparison between the full KV cache (left axis, GB) and index overhead of LycheeCluster (right axis, MB). The percentages annotate the relative size of the index compared to the full KV cache, showing a consistent $\sim 1\%$ overhead.

integrity through terminal punctuation and single newlines. Level 3 addresses intra-sentence pauses using phrasal punctuation, while Level 4 serves as a fine-grained fallback based on whitespace. This prioritization ensures that the text is always split at the most logically significant points available before resorting to finer granularity. While our method utilizes natural language separators, it is designed to be domain-agnostic and robust. First, our hierarchy explicitly prioritizes structural markers common in technical data, such as Markdown code fences and JSON delimiters (Level 1). Second, the chunking algorithm employs a greedy accumulation strategy with a minimum size constraint, ensuring that code snippets with frequent punctuation are not over-segmented but grouped into coherent logical units (e.g., function bodies). Finally, for unstructured or minified inputs where separators are absent, the method degrades to fixed-size chunking, ensuring it performs no worse than traditional baselines (e.g., Quest (Tang et al., 2024)) even in adversarial scenarios. Additionally, while SentenceKV (Zhu et al., 2025) focuses on sentence-level retrieval, its reliance on explicit punctuation limits its applicability to structured data (e.g., code, JSON), which is a core capability of our domain-agnostic LycheeCluster. Therefore, we prioritize comparisons with general-purpose methods like Quest and ClusterKV.

C Memory Overhead Analysis of the Indexing Structure

A critical concern regarding indexing-based sparse attention is whether the auxiliary storage required

Priority Level	Category	Details & Examples
Level-1	Structural Separators	Paragraphs ($\backslash n \backslash n$), Markdown ($-$, $***$, $" "$), Structural Language ($\}$, $\]$, $>$)
Level-2	Grammatical & Semantic	Sentence Terminators ($.$, $?$, $!$, \circ , $?$, $!$), Single Newline ($\backslash n$)
Level-3	Phrasal Separators	English Punctuation ($,$, $;$, $:$), Chinese Punctuation ($,$, $;$, $:$, \cdot)
Level-4	Whitespace	Spaces, Tabs

Table 4: Hierarchical classification of text separators for recursive chunking, ordered by semantic priority from structural delimiters to whitespace.

for the index structure itself becomes a bottleneck in ultra-long context scenarios (e.g., 1M tokens). If the index grows disproportionately, it could potentially negate the memory savings achieved by KV cache compression, particularly on memory-constrained edge devices. To address this, we evaluated the memory footprint of the LycheeCluster index compared to the full KV cache on the Llama3.1-8B-Instruct model across increasing context lengths. As shown in Figure 8, the index overhead is negligible.

- **Minimal Storage Footprint:** While the Full KV Cache (blue line, left axis) rapidly consumes gigabytes of memory as context length increases, the index overhead (red line, right axis) remains in the megabyte range.
- **Consistent $\sim 1\%$ Ratio:** Crucially, the ratio of index size to full KV cache size remains stable at approximately 1% (ranging from 1.0% to 1.3%) across all tested lengths up to 64K.

This linear scaling behavior implies that the index structure remains lightweight relative to the massive KV tensors. Consequently, the memory cost of the index is orders of magnitude smaller than the memory reclaimed through sparsification, ensuring that our method remains highly efficient even on hardware with strict memory limitations.

D Stability Analysis in Ultra-Long Generation

To address concerns regarding index quality decay during extended generation scenarios, we conducted a dedicated stability analysis. We define S_t as the set of active clusters (or indices) retrieved at decoding step t . To quantify the consistency of

the retrieval process over time, we introduce two stability metrics:

- **Jaccard Similarity:** This metric measures the immediate temporal consistency by calculating the overlap between cluster sets at adjacent steps. A higher value indicates a stable focus without abrupt context shifting.

$$J(S_t, S_{t-1}) = \frac{|S_t \cap S_{t-1}|}{|S_t \cup S_{t-1}|} \quad (3)$$

- **Window Hit Rate:** This metric assesses the locality of the retrieval by calculating the proportion of currently retrieved clusters that have appeared within a recent history window w (where $w = 32$). High values imply that the model consistently attends to a relevant working set of information.

$$\text{WindowHit}(S_t) = \frac{|S_t \cap (\bigcup_{i=1}^w S_{t-i})|}{|S_t|} \quad (4)$$

As shown in Figure 9, our method demonstrates strong robustness throughout the generation process. First, the window hit rate (green line) consistently remains close to 1.0 throughout the generation trajectory. This indicates that the model maintains a highly coherent working memory, effectively focusing on relevant historical context without experiencing context loss or generating irrelevant content, even with extremely long sequence lengths. Second, the Jaccard similarity (blue line) maintains a high baseline level. Although there are some fluctuations, especially after 6k-8k steps, this does not represent model failure. On the contrary, it reflects the dynamic adaptability of chain-of-thought generation: as the generation topic continuously evolves in long sequences, the retrieved clusters naturally change to adapt to the new semantic

content. Importantly, LycheeCluster avoids catastrophic collapse, maintaining effective retrieval capabilities even in the later stages of generation where previous methods typically experience degradation. Considering that drift of cluster centroids, future work could also perform global re-clustering when necessary (Sun et al., 2025b).

E Sensitivity Analysis on Clustering Granularity

A critical hyperparameter in LycheeCluster is the granularity of the hierarchical index, specifically the number of fine-grained clusters L . Since L naturally scales with the context length, we analyze the impact of the *average cluster size* (i.e., the average number of chunks assigned to each fine cluster) on both retrieval performance and system efficiency. We conducted this sensitivity analysis on the gov_report subset of LongBench V2. Figure 10 illustrates the trade-off between the recall rate and prefill latency as the average cluster size increases from 1 to 8. The results highlight a clear tension between semantic precision and construction cost: (1) When the average cluster size is small (e.g., 1 chunk per cluster), the centroid provides a highly accurate representation of the underlying data, resulting in the highest recall ($\sim 50\%$). However, as we aggregate more chunks into a single cluster (increasing the x-axis), the centroid becomes a coarser approximation of its constituents. This increased variance within clusters leads to a monotonic decrease in recall, dropping below 40% when averaging 8 chunks per cluster. (2) Conversely, the latency of the prefill phase—dominated by the k -means clustering overhead—is inversely proportional to the cluster size. Smaller cluster sizes imply a larger number of centroids (L) to compute and update, resulting in higher latency (~ 0.68 s at size 1). Increasing the chunks per cluster reduces L , significantly accelerating the index construction process.

Based on these observations, an average of 2 chunks per cluster seems to be the optimal setting for our experiments. This configuration strikes a balance, achieving a significant reduction in prefill latency compared to the baseline (size 1) while sacrificing minimal retrieval accuracy ($\sim 3\%$ drop in recall), whereas larger cluster sizes result in unacceptable performance degradation.

F Extended Analysis and Discussions

In this section, we address specific behaviors of LycheeCluster regarding budget sensitivity and provide a theoretical discussion on the time complexity of our hierarchical retrieval.

F.1 Behavior in Budget-Sufficient Scenarios

We analyze the behavior of LycheeCluster when the context length does not strictly exceed the retrieval budget, particularly in reasoning-intensive tasks. A critical design principle of LycheeCluster is its adaptive nature: the sparse retrieval mechanism is only triggered when the total sequence length (input context + generated chain-of-thought) exceeds the pre-defined token budget. In scenarios where the sum of the input prompt and the generated chain-of-thought remains within the budget, LycheeCluster (and ClusterKV) naturally degenerates to full attention. No tokens are pruned, and no approximation error is introduced. This explains the exclusion of ClusterKV in Table 2 comparisons. Since the average context length in MATH-500 is relatively short, simple paging or clustering methods often operate in a regime identical to full attention unless an extremely small budget is forced. Consequently, we assert that in short-context tasks without budget constraints, LycheeCluster incurs zero degradation compared to full attention. It does not suffer from the unnecessary precision loss that affects methods which enforce static sparsity patterns regardless of buffer availability. However, once the generation extends beyond budget, LycheeCluster seamlessly transitions to hierarchical retrieval, preserving the critical semantic chunks while strictly adhering to the memory bound.

F.2 Complexity Analysis: Sub-linear vs. Logarithmic

A key question regarding hierarchical indexing is the theoretical bound of the retrieval complexity.

Currently, LycheeCluster employs a fixed three-tier structure (coarse unit \rightarrow fine cluster \rightarrow chunk). This design yields a sub-linear time complexity, approximately $O(\sqrt{N})$ depending on the branching factor configuration, rather than strict logarithmic complexity $O(\log N)$.

Why not dynamic depth for $O(\log N)$? While it is theoretically feasible to implement a dynamic tree where the height grows with context length

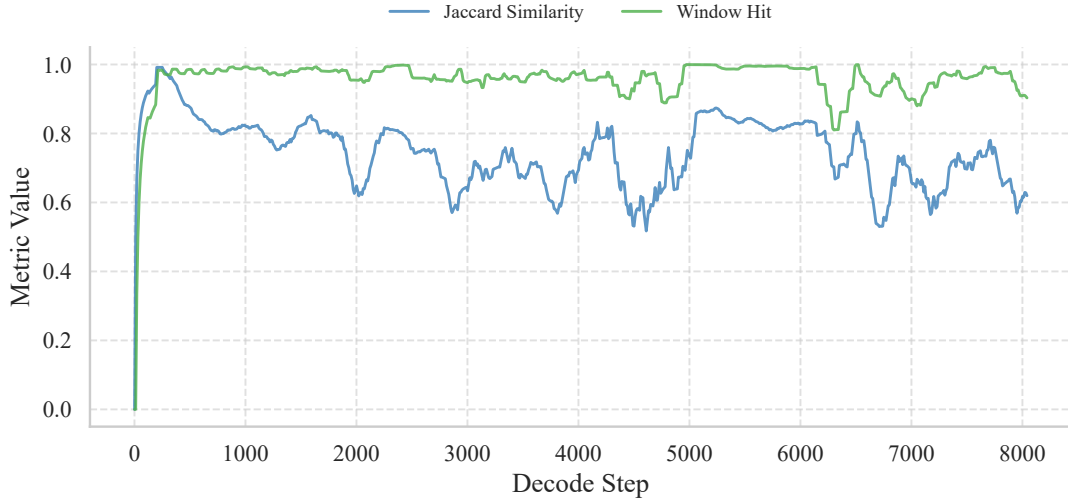


Figure 9: Stability analysis during long-context generation. The blue represents the step-to-step Jaccard Similarity, and the green represents the Window Hit rate ($w = 32$). Stability remains robust initially but shows signs of decay after 6k decode steps.

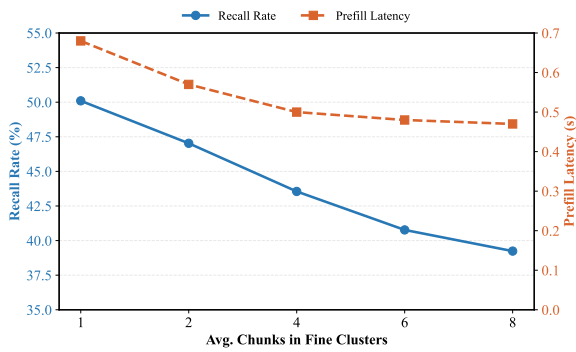


Figure 10: Sensitivity analysis on clustering granularity. We investigate the impact of the average number of chunks per fine cluster on retrieval Recall Rate (blue, left axis) and Prefill Latency (orange, right axis) using the `gov_report` subset.

(e.g., a B-tree or hierarchical k -means with dynamic depth) to achieve $O(\log N)$ retrieval, we chose a fixed-depth approach for practical efficiency:

- **Overhead Trade-off:** For context windows currently relevant to LLMs (e.g., up to 1M tokens), the constant factors and pointer-chasing overheads introduced by maintaining and traversing a deep, dynamic tree outweigh the theoretical gains of logarithmic complexity.
- **Parallelism:** A flatter, fixed-depth hierarchy maps more efficiently to GPU parallelization (e.g., computing distances to all coarse centroids in parallel) compared to the sequential dependency inherent in traversing a deep tree.

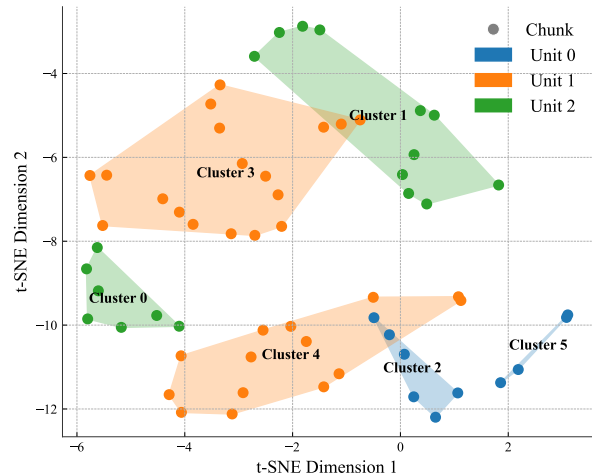


Figure 11: Each point represents a chunk. Colors denote coarse units, while shaded polygons indicate fine clusters.

Therefore, our current design represents an engineering optimum prioritizing wall-clock speedup over asymptotic optimality.

F.3 Visualization of Hierarchical KV Indexing

To intuitively understand how LycheeCluster organizes the KV cache in the semantic space, we employ t-Distributed Stochastic Neighbor Embedding (t-SNE) to visualize the hierarchical structure of the indexed chunks. We randomly sampled a sequence from the *Long In-context Learning* task of LongBench V2, extracting the representative key vectors \bar{k}_j of all chunks constructed during the prefill phase.

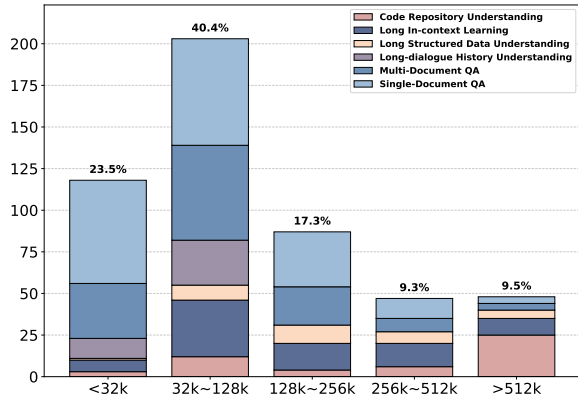


Figure 12: Statistical distribution of data length ranges in the LONGBENCH V2 dataset, using the Llama3.1-8B-instruct tokenizer.

As illustrated in Figure 11, we project these high-dimensional vectors into a 2D space. In this visualization:

- Each **point** represents an atomic chunk (s_j).
- The **shaded polygons** encompass chunks belonging to the same *fine cluster* (c_j).
- The **colors** distinguish different *Coarse Units* (g_p , e.g., Unit 0 in blue, Unit 1 in orange).

The visualization reveals a clear and well-structured semantic topology. First, chunks within the same fine cluster (e.g., Cluster 5 in Unit 0) are tightly grouped, validating our assumption that structure-aware chunks preserve local semantic coherence. Second, the fine clusters themselves are spatially aggregated into distinct coarse units. For instance, the orange region (Unit 1) effectively covers a broad semantic area containing Clusters 3 and 4, while being clearly separated from the green region (Unit 2).

This strong spatial locality confirms the effectiveness of our pruning strategy: a query vector falling near the orange centroid can safely discard the blue and green branches, as their constituent chunks lie far outside the relevant semantic radius. This hierarchical separation is the cornerstone of LycheeCluster’s ability to maintain high recall with sub-linear retrieval complexity.

G Dataset

G.1 LongBench V2

LONGBENCH V2 (Bai et al., 2025) is a benchmark designed to evaluate the deep understanding and reasoning capabilities of LLMs on long

contexts in real-world multi-task scenarios. The dataset consists of 503 highly challenging multiple-choice questions with context lengths ranging from 8k to 2M words (most are below 128k), covering six major task categories: single-document question answering, multi-document question answering, long-context learning, long dialogue history understanding, code repository understanding, and long structured data understanding. To ensure the breadth and practicality of the data, the questions were constructed by nearly a hundred highly educated individuals with diverse professional backgrounds and underwent rigorous review. The difficulty is extremely high; even human experts using search tools only achieved a 53.7% accuracy rate within 15 minutes. Tables 5 show the specific statistical information for the six task categories, and Figure 12 shows the distribution of the dataset across different length ranges. As shown in the figures, LONGBENCH V2 is extremely challenging, with its core evaluation range concentrated in the medium-to-long text range of 32k to 128k (40.4%), while also retaining a considerable proportion of ultra-long text (>512k) test cases, with a maximum text length of up to 4M.

To provide a comprehensive assessment of model performance under varying memory and computational constraints, LONGBENCH V2 categorizes its test instances into three distinct length groups based on the input context size: *Short* (0–32k tokens), *Medium* (32k–128k tokens), and *Long* (exceeding 128k tokens, up to 2M). The *Overall* score represents the macro-average across all instances. This stratified evaluation is crucial for identifying the performance degradation of LLMs, as many models exhibit significant degradation when the context exceeds their effective attention window or fine-tuned limits. By isolating these categories, we can distinguish between the basic reasoning capabilities and its robustness in handling extreme-scale information retrieval and synthesis.

G.2 MATH500

To evaluate the advanced mathematical reasoning capabilities of LLMs, we utilize the MATH500 benchmark (Lightman et al., 2023). This dataset is a high-quality subset of the original MATH benchmark (Hendrycks et al., 2021b), consisting of 500 problems curated to represent a broad spectrum of difficulty levels and mathematical topics. The problems are sourced from elite mathematics competitions, including the American Mathematics Com-

Task	# Samples	Avg. Length	Max Length	Expert Acc.	Expert Time
Code Repository Understanding	50	1M	4M	44%	6.4 min
Long In-context Learning	81	259k	1.5M	63%	8.3 min
Long Structured Data Understanding	33	390k	3M	73%	6.4 min
Long-dialogue History Understanding	39	74k	120k	79%	8.2 min
Multi-Document QA	125	127k	1.7M	36%	6.1 min
Single-Document QA	175	110k	860k	55%	8.9 min

Table 5: Data statistics and expert performance across tasks of LONGBENCH V2.

Context / Task	single	multikey	multivalue	multiquery	vt	fwe	qa1	qa2	Avg.
Full Attention									
4k	100.0	100.0	100.0	100.0	82.5	84.4	89.8	59.2	89.5
8k	91.8	100.0	100.0	99.0	79.6	87.1	81.6	59.2	87.3
16k	63.3	100.0	98.5	98.5	82.0	93.9	79.6	51.0	83.3
32k	95.9	95.9	97.5	97.5	76.3	76.9	79.6	59.2	84.8
LycheeCluster									
4k	100.0	100.0	100.0	99.5	81.2	80.3	89.8	59.2	88.8
8k	100.0	100.0	96.4	100.0	78.0	83.7	81.6	59.2	87.4
16k	93.9	100.0	94.9	94.9	73.1	89.8	77.6	51.0	84.4
32k	95.9	98.0	92.9	93.9	85.7	76.2	75.5	59.2	84.7

Table 6: Performance comparison on RULER benchmark. Note that LycheeCluster achieves comparable or even superior average performance compared to Full Attention, especially at 16k length.

petitions (AMC 10/12) and the American Invitational Mathematics Examination (AIME). MATH-500 spans seven core subjects: Algebra, Geometry, Number Theory, Counting & Probability, Prealgebra, Intermediate Algebra, and Precalculus. Due to its requirement for multi-step logical deduction and precise symbolic manipulation, MATH-500 has become a standard metric for assessing the chain-of-thought reasoning performance of state-of-the-art models.

H More Experiment Results on RULER Benchmark

To assess the long-context capabilities, we evaluate our method on the RULER benchmark (Hsieh et al., 2024). RULER serves as a comprehensive testbed for long-context language models, expanding beyond simple retrieval to include complex aggregation and reasoning tasks. We report results on a diverse set of tasks: *single*, *multikey*, *multivalue*, *multiquery*, *vt*, *fwe*, *qa1*, and *qa2*. We compare **LycheeCluster** against the **Full Attention** baseline across sequence lengths from 4k to 32k.

The experimental results are presented in Table 6. Our method demonstrates remarkable stability and robustness across all context lengths. In the 4k setting, LycheeCluster achieves an average score of 88.8, comparable to the Full Attention baseline. Notably, as the sequence length extends,

LycheeCluster exhibits superior stability in specific scenarios. At 16k length, our method outperforms the full attention model, primarily because the full attention baseline suffers a significant performance drop in the *single* task (dropping to 63.3), whereas LycheeCluster maintains a high score of 93.9. Even at the extreme length of 32k, our approach maintains parity with the full attention mechanism, proving that LycheeCluster can effectively handle extensive contexts without compromising performance.