SemShareKV: Efficient KVCache Sharing for Semantically Similar Prompts via Token-Level LSH Matching

Xinye Zhao
University of Notre Dame
xzhao24@nd.edu

Spyridon Mastorakis
University of Notre Dame
smastor2@nd.edu

Abstract

As large language models (LLMs) continue to scale, the memory footprint of Key-Value (KV) caches during inference has become a significant bottleneck. Existing approaches primarily focus on compressing KV caches within a single prompt or reusing shared prefixes or frequently occurred text segments across prompts. However, such strategies are limited in scenarios where prompts are semantically similar but lexically different, which frequently occurs in tasks such as multi-document summarization and conversational agents. We propose *SemShareKV*, a KV cache sharing and compression framework that accelerates LLM inference by reusing KV cache in semantically similar prompts. Instead of relying on exact token matches, SemShareKV applies fuzzy token matching using Locality-Sensitive Hashing (LSH) on token embeddings and incorporates Rotary Position Embedding (RoPE) to better preserve positional information. By selectively reusing relevant KV pairs from a reference prompt's cache, SemShareKV reduces redundant computation while maintaining output quality. Experiments on diverse summarization datasets show up to $6.25\times$ speedup and 42% lower GPU memory usage with 5k tokens input, with negligible quality degradation. These results highlight the potential of semantic-aware cache sharing for efficient LLM inference.

1 Introduction

Large Language Models (LLMs) have exhibited a strong capability to understand and process human languages, and have been shown to perform comparably to humans in several fields, such as math inference, text memorization, information extraction, story telling (Naveed et al., 2023). Recently released LLMs have significantly advanced in processing and comprehending extremely long prompts. However, this introduces a notable challenge: increased computational demand due to the quadratic time complexity of their Decoder-Only Transformer architecture when handling lengthy text sequences. The issue is further compounded during inference, as the auto-regressive decoding process repeats the computation for each newly generated token (Luohe et al., 2024).

Existing KV cache optimization approaches primarily focus on single-prompt compression through various techniques: (Yang et al., 2024a) leverage decaying KV importance across layers for selective extraction (though with limited small-batch gains), (Gim et al., 2024) employ restrictive markup schemas for text chunk reuse, and (Yao et al., 2024) propose deviation-based recomputation that requires impractical per-chunk precomputation for long inputs. Crucially, these methods operate within the constrained paradigm of single-prompt optimization, failing to exploit the substantial efficiency potential of cross-prompt cache reuse, a significant oversight given the prevalence of semantically similar queries in real-world applications where shared computational savings could be substantial.

Motivated by this challenge, we aim to address the following research question: ***Can we reuse the precomputed KV cache for prompts that are semantically similar?***

To answer this question, we proposed ***SemShareKV***, a KV cache framework that can reuse the cache from one prompt for another that is semantically similar to each other via fuzzy token match. It speeds up prefill phase and compress KV cache in memory. We show that our method can reduce the pre-fill phase time by $6.25\times$ and save 42% GPU memory space. We make the following contributions.

- We introduce SemShareKV, which explores KV cache sharing across semantically similar prompts based on fuzzy token match.

- We evaluate SemShareKV across multiple

datasets, demonstrating its effectiveness in accelerating the prefill phase while simultaneously reducing KV cache size.

- We explored the role of position encoding in KV cache by injecting it into vector embeddings.

2 Related Work

Prior research on KV cache optimization can be categorized into three key directions: (i) **Conventional KVCache Compression**, which focuses on reducing the storage and computational overhead of KV cache by applying quantization, pruning, or other compression techniques; (ii) **KVCache Sharing**, which explores methods to reuse KV cache across different queries or tasks to improve efficiency while maintaining response quality; and (iii) **KVCache Reusing**, which investigates strategies to adapt and re-purpose precomputed KV cache for semantically similar inputs, minimizing redundant computation while preserving model accuracy.

2.1 Conventional KVCache Compression

To address long-context processing, many works propose optimizing inference by retaining only informative tokens. Token-level compression often uses attention-based token selection (Zhang et al., 2023; Xiao et al., 2024; Li et al., 2024; Yang et al., 2024a; Zhong et al., 2024), low-rank decomposition (Sun et al., 2024), or quantization (Zhang et al., 2024; Wang et al., 2024). Model-level approaches redesign architectures to improve reuse (Sun et al., 2025; Yan et al., 2024), while system-level methods focus on memory and scheduling (Kwon et al., 2023; Sheng et al., 2023). Recent work has highlighted the use of value vectors to facilitate compression (Guo et al., 2024).

2.2 KVCache Sharing

Some also emphasize reusing portions of the cache for future or similar queries and prompts. For example, PromptCache (Gim et al., 2024) stores text segments that appear frequently on an inference server using a schema, although this approach hampers usability, as users must conform their natural language to the schema format. Mooncake (Qin et al., 2024), KVSharer (Yang et al., 2024b) and MiniCache (Liu et al.) exploit the high similarity of attention scores among adjacent transformer layers to improve KV cache reuse. By consolidating

or sharing KV pairs between similar layers, these methods improve memory efficiency and streamline token processing. However, their approaches are restricted to sharing in the layer or text segment within adjacent layers or the same LLM, limiting the broader applicability; GPTCache (Regmi and Pun, 2024), (Rasool et al., 2024) and (Bang, 2023) utilize similarity search among queries to reuse KV cache. However, they have a high probability of missing a hit and require the entire query to be similar, offering limited flexibility.

2.3 KVCache Reusing

Limited attention has been directed toward the sharing of KV cache in LLMs. DroidSpeak (Liu et al., 2024b) improves context sharing between fine-tuned LLMs by identifying critical KV cache layers and selectively recomputing them for efficient reuse while maintaining accuracy. LM-Cache (Cheng et al., 2024) introduces a Knowledge Delivery Network (KDN) to optimize KV cache storage and transfer, allowing cost-effective knowledge injection in LLM inference. CacheBlend (Yao et al., 2024), KVShare (Yang et al., 2025), and EPIC (Hu et al., 2024) rely on exact context matching, which is unsuitable for real user scenarios. SentenceKV (Zhu et al., 2025) suffers from inter-sentence information loss, as noted in the CacheBlend. In contrast, SemShareKV introduces RoPE in token matching to address this issue.

3 Observations and Insights

We present three key insights derived from our experiments on three LLMs: Mistral-7B (Jiang et al., 2023), LLaMA-3.1-8B (Grattafiori et al., 2024), and MPT-7B (Team, 2023). These insights show consistent patterns across different LLMs, supporting the generality of our observations.

Insights 1 *HD tokens stay consistent across layers.*

When reusing KV caches from semantically similar prompts, we ensure the reused cache maintains high fidelity with fully recomputed caches to prevent performance degradation. To compare the similarity between two KV matrices, we used our augmented MultiNews (Bai et al., 2023) dataset, where each sample consists of a pair of semantically similar prompts: the ***Target Prompt***, which serves as the primary input to the model, and the ***Reference Prompt*** (Ref Prompt), which acts as the semantically similar counterpart. For each of the
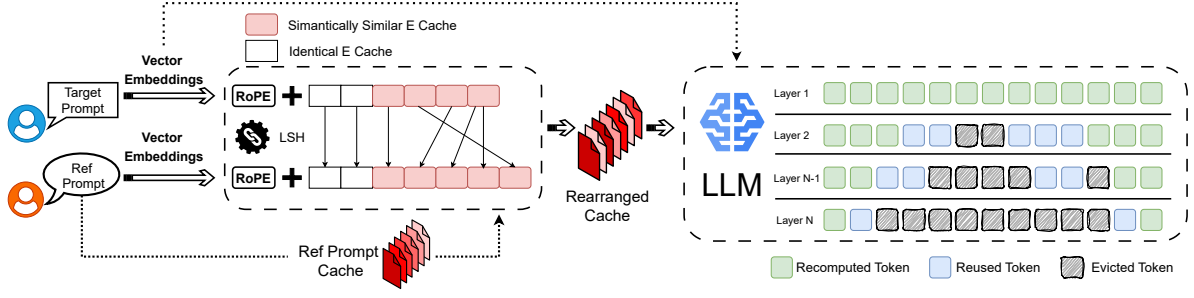
Figure 1: Schematic Overview of SemshareKV

Figure 2: Insight 1: High-deviation tokens remain consistent across layers.

Figure 3: Insight 2: Deeper layers attend to fewer tokens.

aforementioned LLMs, we first computed the KV caches for the prompt pairs independently. Subsequently, we calculated the deviations between the KV caches of the target and reference prompts using the previously mentioned L_2 norm. Tokens with the highest 40% deviation were identified as **High Deviation** (HD) tokens. To further quantify this observation, we computed the Spearman correlation of HD tokens between adjacent layers. As shown in Figure 2, adjacent layers exhibit relatively high consistency in HD token positions.

Insights 2 *Deeper layers focuses on fewer tokens*

To analyze attention patterns across layers, we first averaged the attention scores across all heads in each layer and then computed the mean along the first dimension, resulting in a one-dimensional vector per layer. To quantify this behavior, we introduce **Attention Recovery** (AR), defined as follows:

$$AR = \min\{k \in [n] \mid \frac{\sum_{i=1}^{k} T_i}{\sum_{i=1}^{n} T_i} \geq \text{Thres}\} \quad (1)$$

Where T is a sorted vector of average attention scores for each token, S_{total} represents the total attention score derived from the averaged self-attention matrices, and $Thres$ indicates the threshold of attention score. AR indicates the number of

tokens that must be summed from highest to lowest based on their average attention scores in order to cover $Thres\%$ of the total attention score. We computed AR for each layer, and the results (Figure 3) reveal a consistent trend: as depth increases, AR decreases across all three LLMs, despite minor fluctuations. This suggests that deeper layers concentrate attention on progressively fewer tokens, reflecting more selective focus.

Insights 3 *Deeper layers have more redundant information.*

To reduce memory overhead from the KV cache, a key optimization strategy is to remove tokens containing redundant information. Such tokens contribute minimally to the prediction of next tokens during decoding but occupy substantial GPU memory. However, selective token retention risks information loss, necessitating careful trade-offs between memory savings and generation quality. We evaluate three KV cache retention strategies using perplexity: Constant, with equal retention across layers; Exponential Growth, with higher retention in shallow layers; and Exponential Decay, with more retained in deeper layers. More details of the three retention patterns can be found in Figure C2 in the Appendix.

We applied these three retention patterns to

(a) Llama3.1-8B Retention Pattern   (b) Mistral-7B Retention Pattern   (c) MPT-7B Retention Pattern

Figure 4: Insight 3: Deeper layers contain more redundant information.

LLMs and evaluated generation performance using perplexity. Across all three models, the Exponential Decay pattern achieves the lowest perplexity, indicating the best performance and suggesting this pattern aligns with how LLMs interpret prompt knowledge. Notably, MPT-7B exhibits a spike in AR at the final layer due to its structural differences from Mistral and Llama3.2: unlike the RoPE-based models, MPT-7B employs ALiBi positional encoding and a pre-norm design. ALiBi biases intermediate layers toward recent tokens, while the final layer compensates by attending to more distant tokens, resulting in broader attention patterns and higher AR values.

4 Methodology

4.1 Relevant Concepts

Our work focuses on three critical cache components in modern LLMs:

- **Key Cache (K):** Key vectors encode the structural relationships among tokens in a sequence.

- **Value Cache (V):** Value vectors containing the actual content representations aggregated through attention weights. These preserve the contextual information of each token.

- **Embedding Cache (E):** Contextualized embeddings capturing fundamental semantic and syntactic relationships (Mikolov et al., 2013), providing the fundamental token representations before transformer processing.

4.2 Model Overview

The design of SemShareKV, illustrated in Figure 1, is based on three key insights from Section 3. Our approach employs two core strategies:

- **Recomputation Strategy (Insights 1 & 2):** Prioritize the recomputing of more tokens in shallow layers while reducing the recomputation in

deeper layers, reflecting the varying importance of the layer depth in attention mechanisms.

- **Retention Strategy (Insights 1 & 3):** Preserve more tokens in shallow layers while evicting tokens from deeper layers, optimizing memory usage without significant accuracy degradation.

SemShareKV stores received prompts and their corresponding contextualized E caches in CPU memory. When the LLM receives a new prompt as the target prompt, it retrieves a reference prompt by computing an LSH-distance-based similarity score between the target's contextualized E cache and all stored E caches. The reference prompt with the highest similarity is then loaded onto the GPU along with its corresponding KV cache for reuse.

Once a reference prompt is retrieved, SemShareKV first applies RoPE to the E caches of both the target and reference prompts. Then it uses LSH to match each token in the target prompt to its most similar tokens in the reference prompt. Based on these LSH mappings, the precomputed KV cache of the reference prompt is rearranged token by token and injected into LLM transformer layers. On the first transformer layer, all tokens undergo full recomputation. The recomputed outputs are compared with the rearranged cache values via L_2 norm, identifying high-deviation tokens for prioritized recomputation in subsequent layers. Simultaneously, the system evicts tokens with the lowest attention scores from recent computations, optimizing KV cache memory usage dynamically.

4.3 KVCache Sharing Challenge

The primary challenge in cross-prompt KV cache sharing stems from length disparity between prompts. Inspired by (Liu et al., 2024b), we incorporates positional encoding within the E Cache to enable accurate token alignment while preserving contextual relationships.

Specifically, we use LSH to identify, for each token in the target prompt, the most similar token in the reference prompt based on their vector representations. We use LSH for efficient token similarity search. This process allowed us to reorder the KV cache of the reference prompt to align with the token sequence of the target prompt. Consequently, the reordered KV cache matches the target prompt's length, with its key-value pairs entirely derived from the original KV cache of the reference prompt. The LLM uses the reordered KV cache to the target prompt. Additional details on LSH are provided in Appendix A.2.

4.3.1 Use Relative Position Encoding to Facilitate Fuzzy Token Match

A fundamental limitation of naive matching using the E cache arises from the absence of positional context in its representation. Since raw vector embeddings lack inherent positional information, LSH fails to maintain crucial sequential relationships when identifying reference-target token correspondences. This positional agnosticism in the E cache consequently produces semantically inferior mapping results.

To address this, we introduce positional encoding into the E cache to enhance fuzzy token matching. Two widely used positional encoding strategies are absolute positional encoding (Vaswani et al., 2017), which embeds explicit position information, and relative positional encoding (Su et al., 2024), which captures positional relationships between tokens. In our work, we incorporate Rotary Position Embedding (RoPE) into the E cache and evaluate its impact. Specifically, RoPE is applied to the non-contextual embeddings (E cache) of both the reference and target prompts. Then, LSH is used to match each token in the target prompt's E cache with the most similar token in the reference prompt's E cache. This step is crucial because RoPE introduces position-sensitive information, allowing the same token at different positions to carry distinct semantics, enabling LSH to achieve more accurate token-level matching.

Figure 6 further illustrates how maintaining positional relationships through RoPE improves alignment accuracy, leading to better token retrieval and overall performance. Figure 6a compares the token positions in the original E cache with the rearranged positions after passing through the fuzzy token matching block, while Figure 6b presents the results when using an E cache with positional en-


(a) Standard KVCache Store   (b) Modified KVCache Store

Figure 5: Default vs. modified KV cache storage.

(a) Only E Cache

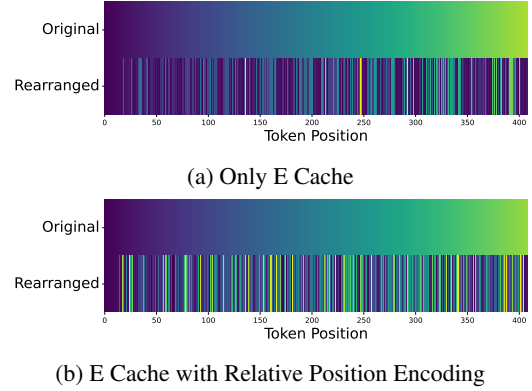(b) E Cache with Relative Position Encoding

Figure 6: Fuzzy matching: with vs. without position encoding.

coding. Notably, the first 20 tokens remain in their original positions, as they represent query tokens. This demonstrates that LSH correctly identifies and preserves query positions.

Beyond the initial tokens, a key difference emerges: without positional encoding, many tokens in the target prompt map to the initial tokens, whereas with positional encoding, they align more accurately with later tokens. We interpret this as a manifestation of "*attention sink*", a phenomenon in self-attention mechanisms where a significant portion of attention scores is consistently assigned to initial tokens, regardless of their actual relevance to the task (Xiao et al., 2023; Fei et al., 2025). Incorporating positional encoding into the E cache effectively mitigates this issue, leading to more accurate token matching and improved performance.

4.3.2 Impact from Position Encoding in KVCache

The second challenge is that LSH-based token rearrangement disrupts position encoding in the precomputed KV cache, affecting the KV matrices from the prefill phase. Previous studies (Gao et al., 2024) have discussed the impact of different position encoding strategies, suggesting that RoPE can be excluded from KV cache by applying it after storage, as illustrated in Figure 5. In evaluating the

Figure 7: Impact of position encoding on E cache and KV cache deviation.

role of RoPE in the E cache, we compared three configurations of KV cache: (i) **With position encoding**, the standard setting where RoPE is applied before storing KV cache; (ii) **Without position encoding**, where RoPE is not applied during storage; and (iii) **Wthout position encoding but reapplied**, where RoPE is omitted during storage but reapplied after LSH-based reordering. Ideally, the rearranged KV cache should closely match the ground-truth KV cache for the target prompt. To quantify the deviation, we compute the L_2 norm between the rearranged and ground-truth caches. As shown in Figure 7, KV cache with position encoding has the lowest deviation, followed by the version without position encoding, while the configuration with reapplied RoPE gives the highest deviation. This highlights the importance of storing KV pairs after RoPE is applied. Ensuring consistent position encoding between the E and KV cache is essential for LLMs to fully leverage them and achieve optimal generation quality.

4.3.3 Recomputation Strategy

We divide the layers into two groups: the first layer and subsequent layers. Given that LLMs tend to focus more on later tokens (Liu et al., 2024a; Yang et al., 2024a), we categorize tokens into Cold (c) and Hot (h) using a dynamic ratio $r_{dynamic}$ from Attention Recovery with a threshold of 55%, meaning $r_{dynamic}$% of tokens with the highest cumulative attention are selected as Hot, and the rest as Cold. The total number of recomputed tokens is defined as Recomputed $= \omega_c \cdot c + \omega_h \cdot h$, where $\omega_c = 0.1$ and $\omega_h = 0.5$ in the SemShareKV.

Starting from the second layer, token selection follows this rule: based on *Insight 1*, the tokens selected in the next layer are derived from those chosen in the previous layer based on a recompute ratio α_{recomp}% of this layer. Based on *Insight*

2, α_{recomp}% in shallow layers will be relatively small while in deeper layers will be relatively large.

4.3.4 Retention Strategy

Similar to token recomputation, we categorize the layers into two groups: the first and the subsequent layers. On the first layer, the retention ratio is determined also by $r_{dynamic}$, follows $Retained = \max(0.8, r_{dynamic})$. And retained tokens are selected based on average attention scores across the last $(1 - r_{dynamic})$% tokens, and only retain the top $r_{dynamic}$% tokens with highest avg attention scores. In detail, the intuition behind selecting retained tokens is as follows: In the first layer, all hot tokens will be retained. Token eviction occurs only among Cold tokens that are not marked as recomputed. The underlying principle is that recomputed tokens provide better representations of the target prompt. If these tokens are evicted, the computational resources and time spent on recomputing them will be wasted. In subsequent layers, based on *Insight 3*, we should retain fewer tokens.

5 Evaluation and Results

5.1 Experiments Setup

We select a diverse set of datasets covering a broad range of tasks. For Q&A, we use **WikiHow** (Koupaee and Wang, 2018) and **Qasper** (Dasigi et al., 2021). For summarization, we include **MultiNews** (Bai et al., 2023) (multi-document), **SAMsum** (Gliwa et al., 2019) (dialogue), and **BookSum**, **PubMed**, and **Big-Patent** (Kwan et al., 2023), which represent narrative, scientific, and patent documents, respectively; all three are single-document summarization. For code completion, we use **LLC** (Guo et al., 2023). For multiple-choice Q&A, we evaluate on **MMLU** (Hendrycks et al., 2021b,a).

We compared SemShareKV against four baselines: (i) **Fully Recompute**: standard inference using the unmodified model from the Transformers library, where the entire prompt is input without any KV cache reuse; (ii) **SnapKV** (Li et al., 2024): a KV cache management method that accelerates the prefill phase by efficient caching but does not compress the KV cache; (iii) **PyramidKV** (Cai et al., 2024): PyramidKV leverages the Pyramidal Information Funneling pattern in LLMs by dynamically adjusting KV cache sizes across layers, allocating more in lower layers and less in higher ones;

Table 1: Performance comparison between SemShareKV and baseline methods

Method	MultiNews	Wikihow	Qasper	SAMSum	PubMed	BookSum	BigPatent	LCC	MMLU
MISTRAL-7B									
Full KV	22.10	20.50	**17.10**	18.79	24.66	22.44	25.47	22.41	34.00
SemShareKV	23.15	19.38	16.52	**21.22**	24.30	22.50	**26.62**	21.55	32.50
SnapKV	23.07	21.32	15.55	20.16	24.58	23.22	25.78	**25.98**	**35.50**
PyramidKV	**23.71**	20.06	16.34	20.59	**25.13**	**23.87**	26.20	15.83	33.00
H2O	23.04	**21.33**	15.88	20.50	23.53	22.77	24.99	16.57	33.00
LLAMA3.1-8B									
Full KV	22.49	19.71	14.21	16.69	24.50	22.65	27.26	19.01	**55.00**
SemShareKV	23.18	20.41	14.41	**18.61**	24.04	21.66	26.71	**21.39**	51.00
SnapKV	**23.84**	**21.65**	14.70	16.07	**24.82**	**22.76**	27.48	19.16	52.50
PyramidKV	23.71	20.69	**15.56**	16.86	24.50	22.72	**27.53**	14.16	52.00
H2O	22.81	20.61	14.44	16.81	24.19	22.08	26.94	21.32	47.50

(iv) **H2O** (Zhang et al., 2023): a dynamic KV cache eviction strategy that compresses KV memory by prioritizing important tokens, but does not optimize the prefill phase. We use a modified H2O compressing 10% of the cache per layer, with SnapKV and H2O as baselines for prefill optimization and KV cache compression. Our framework is based on a simple LRU mechanism to manage precomputed caches, and SemShareKV can be easily integrated with other cache management strategies. The experiments were run on a single A100 GPU with standard attention. The implementation details are in Appendix D.

To the best of our knowledge, no existing dataset benchmarks LLMs on KV cache sharing across semantically similar prompts. To bridge this gap, we constructed evaluation samples by randomly selecting portions of entries from existing datasets and rewriting them using the Llama3 model. Then, these rewritten samples were manually verified to ensure that they remained semantically close to the originals. More details in the data preparation are provided in the Appendix B.

5.2 Benchmarking Evaluation

We argue that using Fuzzy Token Match introduces only a negligible overhead to model inference. Table 1 reports the ROUGE-L scores (Lin, 2004). Benchmarking results show that SemShareKV achieves performance comparable to or better than other baseline methods. Notably, in most of the evaluated datasets, *Fully Recompute* fails to attain the highest performance scores. We attribute this phenomenon to the token eviction mechanisms employed by SemShareKV, SnapKV, PyramidKV and H2O. By selectively retaining only the most seman-

(a) TTFT Comparison

(b) KVCache Size Comparison

Figure 8: Efficiency Evaluation Results.

tically significant tokens for self-attention computation, these methods effectively reduce redundant information in the semantic representation, thereby enhancing the model's generation quality.

5.3 Efficiency Evaluation

We evaluate SemShareKV based on Time To First Token (TTFT) and KV cache GPU KV memory usage, benchmarking it against Fully Recompute, SnapKV, and the unmodified H2O model. Figure 8 demonstrates the efficiency advantages of SemShareKV on the MultiNews dataset, showing
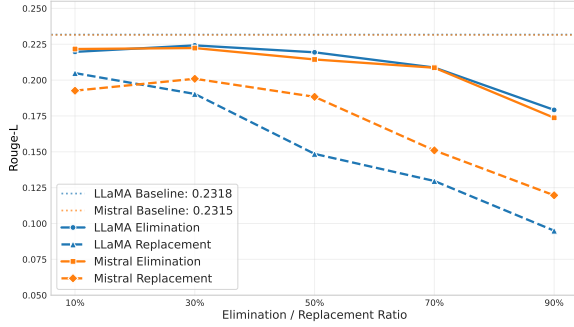
Figure 9: Impact of Prompt Similarity

Figure 10: Cache Retention Ratio and Performance

Table 2: Ablation study on ROUGE-L for SemShareKV and its ablations across datasets.

Method	SAMSum(↑)	MultiNews(↑)
SemShareKV	**21.22**	**23.15**
Fuzzy + Full Cache	17.27	21.34
Ablation-Zero	14.63	17.71
Ablation-Random	5.38	12.67

reasonable even with substantial reductions. Notably, SemShareKV maintains strong performance even when 50% of context sentences are removed, highlighting the effectiveness of LSH-based token-level matching. This trend holds across both LLMs, suggesting the generality of our approach. Additionally, based on the observed performance trend, we empirically set a threshold of 0.8 for applying SemShareKV, meaning that if two prompts have an LSH similarity score above 0.8, SemShareKV can be applied. Figure 9 illustrates how the LSH-distance-based similarity changes as the replacement and elimination ratios increase. More details are in Appendix B.2.

5.5 Ablation Study

We conducted three ablation studies to assess the effect of fuzzy token matching on semantic understanding. First, when applying the fuzzy matching technique with the full cache, we observed a decline in performance, indicating the necessity of the token retention mechanism. Second, when matched KV cache tokens were either zeroed out or replaced with random ones, the ROUGE-L scores dropped significantly compared to the full SemShareKV method, confirming that fuzzy matching is crucial to capture semantic relationships (Table 2). Third, as shown in Figure 10, an analysis of the KV cache compression ratio in the MultiNews dataset revealed that retaining too much cache introduces redundancy and degrades performance, while retaining too little leads to information loss, highlighting the importance of a balanced cache retention strategy.

6 Conclusion

We proposed SemShareKV, a KV cache sharing framework that enables reuse across semantically similar prompts through fuzzy token matching using LSH. SemShareKV achieves a speed of 6.25× and saves up to 42% KV cache memory space compared to conventional KV cache, with a minimum performance drop.

consistent improvements over baseline methods: SemShareKV achieves **6.25×** faster Time-To-First-Token (TTFT) than Fully Recompute, PyramidKV and H2O, **2.23×** faster TTFT than SnapKV, while reducing memory usage by **42%**. Although PyramidKV uses less KV cache than SemShareKV, it does not accelerate the prefill phase, highlighting the effectiveness of SemShareKV. However, as illustrated in Figure 8b, SemShareKV offers limited performance improvements for shorter prompts (fewer than 700 tokens), which we attribute to the overhead caused by fuzzy token matching and the rearrangement of tokens from the precomputed cache of the reference prompt. In future work, our goal is to minimize this overhead.

5.4 Impact of Prompt Similarity on Performance.

To evaluate similarity effects and SemShareKV performance, we designed two studies using the same percentage range (10% to 90% in 10% increments): 1) **Randomly eliminating** sentences from the context, and 2) **Randomly replacing** sentences with others from the MultiNews dataset. We then applied SemShareKV to reuse the cache from the modified reference prompt for the target prompt. As shown in Figure 9, performance gradually degrades as more sentences are removed, yet remains

Limitations

Although SemShareKV effectively shares KV caches across semantically similar prompts, the speedup decreases for shorter prompts because of the overhead introduced by fuzzy token matching, and several hyper-parameter require careful tuning. Our current implementation focuses on demonstrating the effectiveness of SemShareKV and does not yet support FlashAttention, which we plan to explore into future work. The matching threshold is also empirically determined, and exploring adaptive strategies remains an open direction. In the ablation study, we replaced sentences with other samples from the same dataset, ensuring a fair comparison within a consistent data distribution. Although cross-corpus replacement could test broader robustness, our setup effectively isolates the impact of semantic mismatches. This opens future work not only to extend the experiments to cross-corpus settings but also to construct datasets of semantically similar prompts to better evaluate robustness.

References

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. 2023. Longbench: A bilingual, multitask benchmark for long context understanding. *arXiv preprint arXiv:2308.14508*.

Fu Bang. 2023. Gptcache: An open-source semantic cache for llm applications enabling faster answers and cost savings. In *Proceedings of the 3rd Workshop for Natural Language Processing Open Source Software (NLP-OSS 2023)*, pages 212–218.

Zefan Cai, Yichi Zhang, Bofei Gao, Yuliang Liu, Tianyu Liu, Keming Lu, Wayne Xiong, Yue Dong, Baobao Chang, Junjie Hu, and 1 others. 2024. Pyramidkv: Dynamic kv cache compression based on pyramidal information funneling. *arXiv preprint arXiv:2406.02069*.

Yihua Cheng, Kuntai Du, Jiayi Yao, and Junchen Jiang. 2024. Do large language models need a content delivery network? *arXiv preprint arXiv:2409.13761*.

Pradeep Dasigi, Kyle Lo, Iz Beltagy, Arman Cohan, Noah A. Smith, and Matt Gardner. 2021. A dataset of information-seeking questions and answers anchored in research papers.

Mayur Datar, Nicole Immorlica, Piotr Indyk, and Vahab S Mirrokni. 2004. Locality-sensitive hashing scheme based on p-stable distributions. In *Proceedings of the twentieth annual symposium on Computational geometry*, pages 253–262.

Matthijs Douze, Alexandr Guzhva, Chengqi Deng, Jeff Johnson, Gergely Szilvasy, Pierre-Emmanuel Mazaré, Maria Lomeli, Lucas Hosseini, and Hervé Jégou. 2024. The faiss library.

Weizhi Fei, Xueyan Niu, Guoqing Xie, Yingqing Liu, Bo Bai, and Wei Han. 2025. Efficient prompt compression with evaluator heads for long-context transformer inference. *arXiv preprint arXiv:2501.12959*.

Bin Gao, Zhuomin He, Puru Sharma, Qingxuan Kang, Djordje Jevdjic, Junbo Deng, Xingkun Yang, Zhou Yu, and Pengfei Zuo. 2024. Cost-efficient large language model serving for multi-turn conversations with cachedattention. In *USENIX Annual Technical Conference*.

In Gim, Guojun Chen, Seung-seob Lee, Nikhil Sarda, Anurag Khandelwal, and Lin Zhong. 2024. Prompt cache: Modular attention reuse for low-latency inference. *Proceedings of Machine Learning and Systems*, 6:325–338.

Bogdan Gliwa, Iwona Mochol, Maciej Biesek, and Aleksander Wawer. 2019. SAMSum corpus: A human-annotated dialogue dataset for abstractive summarization. In *Proceedings of the 2nd Workshop on New Frontiers in Summarization*, pages 70–79, Hong Kong, China. Association for Computational Linguistics.

Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, and 1 others. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.

Daya Guo, Canwen Xu, Nan Duan, Jian Yin, and Julian McAuley. 2023. Longcoder: A long-range pre-trained language model for code completion. In *International Conference on Machine Learning*, pages 12098–12107. PMLR.

Zhiyu Guo, Hidetaka Kamigaito, and Taro Watanabe. 2024. Attention score is not all you need for token importance indicator in kv cache reduction: Value also matters. *arXiv preprint arXiv:2406.12335*.

Dan Hendrycks, Collin Burns, Steven Basart, Andrew Critch, Jerry Li, Dawn Song, and Jacob Steinhardt. 2021a. Aligning ai with shared human values. *Proceedings of the International Conference on Learning Representations (ICLR)*.

Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2021b. Measuring massive multitask language understanding. *Proceedings of the International Conference on Learning Representations (ICLR)*.

Junhao Hu, Wenrui Huang, Haoyi Wang, Weidong Wang, Tiancheng Hu, Qin Zhang, Hao Feng, Xusheng Chen, Yizhou Shan, and Tao Xie. 2024. Epic: Efficient position-independent context caching for serving large language models. *arXiv preprint arXiv:2410.15332*.

Piotr Indyk and Rajeev Motwani. 1998. Approximate nearest neighbors: towards removing the curse of dimensionality. In *Proceedings of the thirtieth annual ACM symposium on Theory of computing*, pages 604–613.

Albert Q Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, and 1 others. 2023. Mistral 7b. *arXiv preprint arXiv:2310.06825*.

Mahnaz Koupaee and William Yang Wang. 2018. Wikihow: A large scale text summarization dataset. *arXiv preprint arXiv:1810.09305*.

Wai-Chung Kwan, Xingshan Zeng, Yufei Wang, Yusen Sun, Liangyou Li, Lifeng Shang, Qun Liu, and Kam-Fai Wong. 2023. M4LE: A Multi-Ability Multi-Range Multi-Task Multi-Domain Long-Context Evaluation Benchmark for Large Language Models.

Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph E. Gonzalez, Hao Zhang, and Ion Stoica. 2023. Efficient memory management for large language model serving with pagedattention. In *Proceedings of the ACM SIGOPS 29th Symposium on Operating Systems Principles*.

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. *arXiv preprint arXiv:2404.14469*.

Chin-Yew Lin. 2004. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, pages 74–81.

Akide Liu, Jing Liu, Zizheng Pan, Yefei He, Gholamreza Haffari, and Bohan Zhuang. Minicache: Kv cache compression in depth dimension for large language models, 2024b. *URL https://arxiv.org/abs/2405.14366*.

Nelson F Liu, Kevin Lin, John Hewitt, Ashwin Paranjape, Michele Bevilacqua, Fabio Petroni, and Percy Liang. 2024a. Lost in the middle: How language models use long contexts. *Transactions of the Association for Computational Linguistics*, 12:157–173.

Yuhan Liu, Esha Choukse, Shan Lu, Junchen Jiang, and Madan Musuvathi. 2024b. Droidspeak: Enhancing cross-llm communication. *arXiv preprint arXiv:2411.02820*.

Shi Luohe, Zhang Hongyi, Yao Yao, Li Zuchao, and Zhao Hai. 2024. Keep the cost down: A review on methods to optimize llm's kv-cache consumption. *arXiv preprint arXiv:2407.18003*.

Tomas Mikolov, Kai Chen, Gregory S. Corrado, and Jeffrey Dean. 2013. Efficient estimation of word representations in vector space. In *International Conference on Learning Representations*.

Humza Naveed, Asad Ullah Khan, Shi Qiu, Muhammad Saqib, Saeed Anwar, Muhammad Usman, Naveed Akhtar, Nick Barnes, and Ajmal Mian. 2023. A comprehensive overview of large language models. *arXiv preprint arXiv:2307.06435*.

Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, pages 311–318.

Ruoyu Qin, Zheming Li, Weiran He, Mingxing Zhang, Yongwei Wu, Weimin Zheng, and Xinran Xu. 2024. Mooncake: A kvcache-centric disaggregated architecture for llm serving. *arXiv preprint arXiv:2407.00079*.

Zafaryab Rasool, Scott Barnett, David Willie, Stefanus Kurniawan, Sherwin Balugo, Srikanth Thudumu, and Mohamed Abdelrazek. 2024. Llms for test input generation for semantic caches. *arXiv preprint arXiv:2401.08138*.

Sajal Regmi and Chetan Phakami Pun. 2024. Gpt semantic cache: Reducing llm costs and latency via semantic embedding caching. *arXiv preprint arXiv:2411.05276*.

Ying Sheng, Lianmin Zheng, Binhang Yuan, Zhuohan Li, Max Ryabinin, Beidi Chen, Percy Liang, Christopher Ré, Ion Stoica, and Ce Zhang. 2023. Flexgen: High-throughput generative inference of large language models with a single gpu. In *International Conference on Machine Learning*, pages 31094–31116. PMLR.

Jianlin Su, Murtadha Ahmed, Yu Lu, Shengfeng Pan, Wen Bo, and Yunfeng Liu. 2024. Roformer: Enhanced transformer with rotary position embedding. *Neurocomputing*, 568:127063.

Hanshi Sun, Li-Wen Chang, Wenlei Bao, Size Zheng, Ningxin Zheng, Xin Liu, Harry Dong, Yuejie Chi, and Beidi Chen. 2024. Shadowkv: Kv cache in shadows for high-throughput long-context llm inference. *arXiv preprint arXiv:2410.21465*.

Yutao Sun, Li Dong, Yi Zhu, Shaohan Huang, Wenhui Wang, Shuming Ma, Quanlu Zhang, Jianyong Wang, and Furu Wei. 2025. You only cache once: Decoder-decoder architectures for language models. *Advances in Neural Information Processing Systems*, 37:7339–7361.

MosaicML NLP Team. 2023. Introducing mpt-7b: A new standard for open-source, commercially usable llms. Accessed: 2023-05-05.

Ashish Vaswani, Noam M. Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Neural Information Processing Systems*.

Zihao Wang, Bin Cui, and Shaoduo Gan. 2024. Squeezeattention: 2d management of kv-cache in llm inference via layer-wise optimal budget. *arXiv preprint arXiv:2404.04793*.

Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, and 3 others. 2020. Transformers: State-of-the-art natural language processing. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.

Guangxuan Xiao, Jiaming Tang, Jingwei Zuo, Junxian Guo, Shang Yang, Haotian Tang, Yao Fu, and Song Han. 2024. Duoattention: Efficient long-context llm inference with retrieval and streaming heads. *arXiv preprint arXiv:2410.10819*.

Guangxuan Xiao, Yuandong Tian, Beidi Chen, Song Han, and Mike Lewis. 2023. Efficient streaming language models with attention sinks. *arXiv preprint arXiv:2309.17453*.

Ruiqing Yan, Linghan Zheng, Xingbo Du, Han Zou, Yufeng Guo, and Jianfei Yang. 2024. Recurformer: Not all transformer heads need self-attention. *arXiv preprint arXiv:2410.12850*.

Dongjie Yang, XiaoDong Han, Yan Gao, Yao Hu, Shilin Zhang, and Hai Zhao. 2024a. Pyramidinfer: Pyramid kv cache compression for high-throughput llm inference. *arXiv preprint arXiv:2405.12532*.

Huan Yang, Renji Zhang, Mingzhe Huang, Weijun Wang, Yin Tang, Yuanchun Li, Yunxin Liu, and Deyu Zhang. 2025. Kvshare: An llm service system with efficient and effective multi-tenant kv cache reuse. *arXiv preprint arXiv:2503.16525*.

Yifei Yang, Zouying Cao, Qiguang Chen, Libo Qin, Dongjie Yang, Hai Zhao, and Zhi Chen. 2024b. Kvsharer: Efficient inference via layerwise dissimilar kv cache sharing. *arXiv preprint arXiv:2410.18517*.

Jiayi Yao, Hanchen Li, Yuhan Liu, Siddhant Ray, Yihua Cheng, Qizheng Zhang, Kuntai Du, Shan Lu, and Junchen Jiang. 2024. Cacheblend: Fast large language model serving with cached knowledge fusion. *arXiv preprint arXiv:2405.16444*.

Hailin Zhang, Xiaodong Ji, Yilin Chen, Fangcheng Fu, Xupeng Miao, Xiaonan Nie, Weipeng Chen, and Bin Cui. 2024. Pqcache: Product quantization-based kvcache for long context llm inference. *arXiv preprint arXiv:2407.12820*.

Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q Weinberger, and Yoav Artzi. 2019. Bertscore: Evaluating text generation with bert. *arXiv preprint arXiv:1904.09675*.

Zhenyu Zhang, Ying Sheng, Tianyi Zhou, Tianlong Chen, Lianmin Zheng, Ruisi Cai, Zhao Song, Yuandong Tian, Christopher Ré, Clark Barrett, and 1 others. 2023. H2o: Heavy-hitter oracle for efficient generative inference of large language models. *Advances in Neural Information Processing Systems*, 36:34661–34710.

Meizhi Zhong, Xikai Liu, Chen Zhang, Yikun Lei, Yan Gao, Yao Hu, Kehai Chen, and Min Zhang. 2024. Zigzagkv: Dynamic kv cache compression for long-context modeling based on layer uncertainty. *arXiv preprint arXiv:2412.09036*.

Yuxuan Zhu, Ali Falahati, David H. Yang, and Mohammad Mohammadi Amiri. 2025. Sentencekv: Efficient llm inference via sentence-level semantic kv caching. *ArXiv*, abs/2504.00970.

A Formula and Inference

A.1 Rotary Position Encoding

In our methodology, we introduced the application of **Rotary Position Embedding (RoPE)** to the E cache, which improves the performance of fuzzy token matching. RoPE is designed to incorporate positional information directly into embeddings, allowing for improved alignment between tokens in a sequence. This is particularly important in natural language processing tasks where the order of words can significantly impact the meaning and context.

The formula of RoPE in a 2-D case is shown below:

$$\text{RoPE}(\mathbf{x}) = \begin{bmatrix} \cos(\theta_k) & -\sin(\theta_k) \\ \sin(\theta_k) & \cos(\theta_k) \end{bmatrix} \begin{bmatrix} x_{2k} \\ x_{2k+1} \end{bmatrix} \quad (1)$$

In this equation, $\theta_k = 10000^{-2k/d}$, where d represents the embedding dimension. The use of RoPE allows for the effective encoding of relative positional information, enabling the model to better capture the relationships between tokens in a sequence. Integrating RoPE into the E cache facilitates the identification of semantically similar tokens using LSH, leading to more accurate and efficient fuzzy token matching. This enhancement helps the model perform more accurately on tasks that require strong semantic understanding.

A.2 Locality-Sensitive Hashing (LSH)

Locality-Sensitive Hashing (LSH) is a technique that enables efficient approximate nearest neighbor searches in high-dimensional spaces by ensuring similar input items are hashed into the same bucket

Table B1: Similarity evaluation on benchmarking datasets using ROUGE-L, BERTScore, and BLEU.

Metric	MultiNews	Wikihow	Qasper	SAMSum	PubMed	BookSum	BigPatent	LCC	MMLU
N of Samples	100	100	100	100	100	100	100	100	200
Rewrite % (Avg)	54.58	100	28.99	46.75	45.64	44.31	28.55	29.51	76.39
ROUGE-L(%)	84.71	83.82	91.71	50.90	88.15	78.15	90.03	87.34	44.04
BERTScore(%)	95.85	95.98	98.13	86.97	95.48	95.58	96.07	98.41	89.57
BLEU(%)	90.40	87.84	91.32	24.68	89.22	81.16	89.29	89.76	40.51

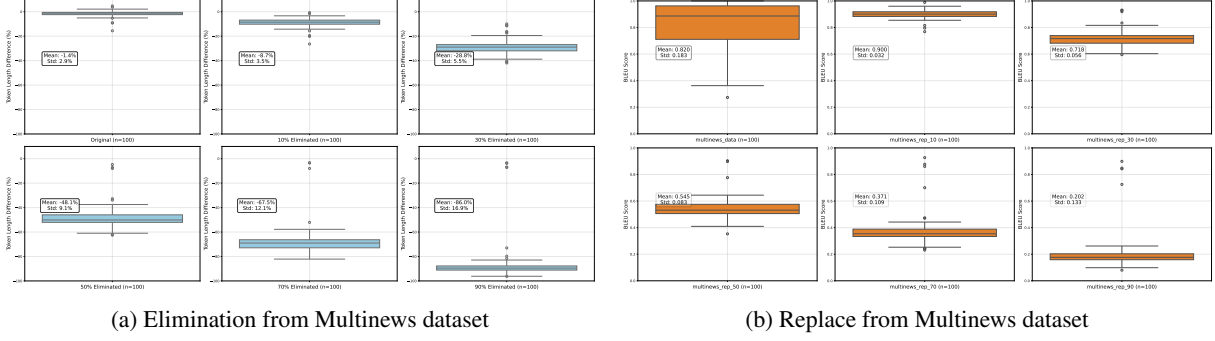(a) Elimination from Multinews dataset (b) Replace from Multinews dataset

Figure C1: Insight 3: Deeper layers contain more redundant information.

with high probability (Indyk and Motwani, 1998). This reduces the number of distance computations required, making LSH particularly useful for large datasets in applications such as image retrieval and natural language processing (Datar et al., 2004). In LSH for Euclidean distance, a common hash function is:

$$h(x) = \lfloor \frac{x \cdot r + b}{w} \rfloor$$

where r is a random vector, b is a random offset, and w is the hash width. This overview encapsulates the theory and practical application of LSH in our framework.

A.3 LSH-Distance Based Similarity Score

For retrieving reference prompts to reuse cache with SEMSHAREKV, we compute a similarity score by normalizing the LSH distance and inverting it to fit within a [0, 1] range:

$$d_{\text{norm}} = \frac{\text{LSH\_dist} - \min(\text{dist})}{\max(\text{dist}) - \min(\text{dist})}$$

$$\text{Similarity} = \text{clip}(1 - d_{\text{norm}}, 0, 1)$$

(2)

where d_{norm} denotes the normalized LSH distance; $\min(\text{dist})$ is set to 0 and $\max(\text{dist})$ is set to 30.

A.4 Key-Value Deviation

We define Key-Value Deviation with L_2 norm as below:

$$\sigma_K = \|K^{\text{reused}} - K^{\text{recomputed}}\|_2,$$
$$\sigma_V = \|V^{\text{reused}} - V^{\text{recomputed}}\|_2,$$
$$\sigma_{KV} = \sigma_K + \sigma_V$$

(3)

Where K^{reused} and V^{reused} represent the Key and Value matrices in cache reused from the semantic similar prompt; $K^{\text{recomputed}}$ and $V^{\text{recomputed}}$ refer to the Key and Value matrices recomputed at the current layer.

A.5 Token Recomputation

The total number of tokens recomputed on layer i is represented as

$$\text{Recomp}[i] = \text{T} \prod_{j=1}^{i} \alpha_{\text{recomp}}[j]$$

(4)

Where T denotes the total number of tokens, i represents the layer index.

A.6 Token Retention

The token retained on each layer is defined as:

$$\text{Retain}[i] = \text{T} \prod_{j=1}^{i} \alpha_{\text{retain}}[j]$$

(5)

Where T is the total tokens, i the layer index, and $\alpha_{\text{retain}}[j]$ the token retention ratio at layer j; tokens not retained are evicted. Typically, α_{retain} is larger in shallow layers and smaller in deeper ones.

B Data Preparation

B.1 Benchmark Datasets

We categorize these nine English-language datasets into four groups based on how semantically similar samples are constructed and the nature of the task.

1. **MultiNews (Bai et al., 2023):** This datasets contain samples composed of multiple independent passages or articles. To generate semantically similar samples, we randomly select one passage or article from each sample and use the Llama 3 model to rewrite it while preserving the original semantics. The rewritten passage is constrained to have a similar length to the original (within a 10% difference in token count). We then replace the original passage with the rewritten one to construct a semantically similar prompt. The position of the rewritten passage naturally varies across samples, appearing at the beginning, middle, or end of the context.

2. **SAMSum (Gliwa et al., 2019), PubMed, BigPatent, BookSum (Kwan et al., 2023), LCC (Guo et al., 2023):** These datasets consist of semantically continuous text or codes. For each sample, we divide the context into individual sentences and randomly select a contiguous segment of the total sentence count. This segment is rewritten using the Llama 3 model, with the constraint that the token count deviates by less than 10% from the original. The rewritten segment replaces the original to create a semantically similar prompt, with its position varying within the context in a similar manner.

3. **Qasper (Dasigi et al., 2021) and Wiki-How (Koupaee and Wang, 2018):** These datasets consist of Q&A tasks where each question must be answered based on a specific provided context. To preserve the accuracy of the questions, we use the LLM to rewrite only part of the context, leaving the questions unchanged.

4. **MMLU (Hendrycks et al., 2021a):** MMLU is a multiple-choice question-answering dataset. To ensure the logical integrity of the questions and preserve the original answers, we prompt the LLM to paraphrase each entire question.

Table B2: Effieicny Comparision with Throughput

| Method	Input Length	Token / Sec)
Fully Recompute	5k	12.681
SemShareKV	5k	34.065
SnapKV	5k	22.576
H2O	5k	16.273
PyramidKV	5k	3.991

Figure C2: The three retention patterns start from the same retention ratio.

Table B1 presents the results of the similarity evaluation, measured using ROUGE-L (Lin, 2004), BLEU (Papineni et al., 2002), and BERTScore (Zhang et al., 2019). We include both longest common subsequence-based metrics (ROUGE-L), n-gram-based metrics (BLEU) and embedding-based metrics (BERTScore) to provide a comprehensive evaluation of semantic similarity across rewritten datasets.

B.2 Eliminination and Replacement Dataset

To study the impact of prompt similarity on LLM performance when applying SemShareKV, we designed two ablation studies. In the first, we randomly removed a portion of sentences from each sample in the MULTINEWS dataset, then applied SemShareKV to evaluate its effectiveness. Figure C1a presents box plots of token length differences in the Elimination datasets compared to the original dataset. Figure C1b shows the BLEU scores of the Replacement datasets relative to the original dataset.

C Extra Experimental Results

Figure C2 illustrates the three retention patterns discussed in **Insights 3**. Table B2 provides additional benchmarking results of Total Batch Time (TBT) and throughput. Since all experiments were conducted with batch size equal to 1, the reported throughput values are equivalent to TBT.

D Implementation and Hyperparameters

SemShareKV is implemented in Python using the `transformers` library (Wolf et al., 2020), with the monkeypatching technique. We use the Locality-Sensitive Hashing from FAISS (Douze et al., 2024) library. The code is available at: `https://github.com/JasperZhao666/SemShareKV-public.git`. Details of the key functions and their roles are outlined below:

- **mistral_attn_forward**: A modified version of `MistralAttention.forward` from the `transformers` library, incorporating the SemShareKV mechanism. The hyperparameters used in our experiments are also specified in this function.

- **replace_mistral_forward**: Applies monkey-patching to substitute the original Mistral model attention forward function in the `transformers` library with our customized SemShareKV implementation.

- **llama_attn_forward**: A modified version of `LlamaAttention.forward` from the `transformers` library, incorporating the SemShareKV mechanism. The hyperparameters used in our experiments are also specified in this function.

- **replace_llama_forward**: Applies monkey-patching to substitute the original Llama model attention forward function in the `transformers` library with our customized SemShareKV implementation.

- **prepare_fuzzy_caches**: Encodes ROPE into E caches and performs fuzzy token matching using locality-sensitive hashing (LSH).

In general, SemShareKV is built on the transformer architecture and consists of fewer than 300 new lines of code, which makes it lightweight and easily transferable to other LLMs.

E Artifact Use and Compliance with Intended Purpose

The datasets used in this study are publicly available and are consistent with their intended use, as specified by the respective sources. In preparing the data, we adhered to ethical guidelines and ensured that the use of these publicly released datasets was for research purposes only.

For the created artifacts, such as the semantically similar samples, we have ensured that the use of these modified datasets remains consistent with the original intended research purpose. The generated data serves the purpose of advancing research in semantic similarity and does not extend beyond the intended scope of the original datasets.