
Fast KVzip: Efficient and Accurate LLM Inference with Gated KV Eviction

Jang-Hyun Kim¹ Dongyoon Han¹ Sangdoo Yun¹

Abstract

Efficient key-value (KV) cache management is crucial for the practical deployment of large language models (LLMs), yet existing compression techniques often incur a trade-off between performance degradation and computational overhead. We propose a novel gating-based KV cache eviction method for frozen-weight LLMs that achieves high compression ratios with negligible computational cost. Our approach introduces lightweight sink-attention gating modules to identify and retain critical KV pairs, and integrates seamlessly into both the prefill and decoding stages. The proposed gate training algorithm relies on forward passes of an LLM, avoiding expensive backpropagation, while achieving strong task generalization through a task-agnostic reconstruction objective. Extensive experiments across the Qwen2.5-1M, Qwen3, and Gemma3 families show that our method maintains near-lossless performance while evicting up to 70% of the KV cache. The results are consistent across a wide range of tasks, including long-context understanding, code comprehension, and mathematical reasoning, demonstrating the generality of our approach.

1. Introduction

Transformer-based LLMs have revolutionized the landscape of artificial intelligence, solving complex challenges ranging from long-context processing to advanced reasoning (Jaech et al., 2024; Comanici et al., 2025). A cornerstone of this success is the attention mechanism, which allows models to effectively contextualize features across sequences (Vaswani et al., 2017). To mitigate the computational cost of re-calculating attention features at every step, inference systems adopt KV caching (Dai et al., 2019).

¹NAVER AI Lab. Correspondence to: Jang-Hyun Kim <kimjanghyun1230@gmail.com>, Sangdoo Yun <sangdoo.yun@navercorp.com>. The source code is available at <https://github.com/Janghyun1230/fastkvzip>.

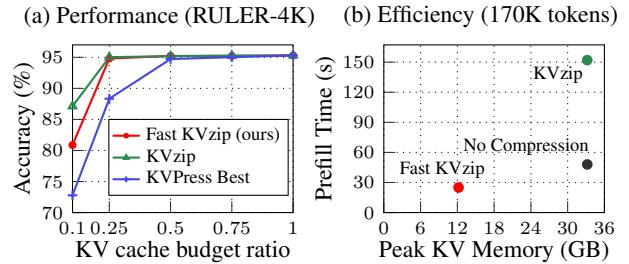


Figure 1. **Highlighted results.** (a) Evaluation on the KVPress benchmark (Devoto et al., 2025) (RULER-4K dataset, Qwen3-8B) compared to state-of-the-art methods from December 2025. (b) Chunked-prefill efficiency for 170K tokens using Qwen2.5-14B-1M with 30% KV cache budget and BF16 precision.

However, the effectiveness of KV caching is increasingly strained by the growing demand for long-context inference. As the memory required to store KV features scales linearly with sequence length, KV caches become a dominant bottleneck during inference (Kwon et al., 2023). While KV cache eviction approaches alleviate this burden, they often face a critical dilemma: Methods with negligible compression overhead result in significant performance degradation (Zhang et al., 2023; Li et al., 2024), while methods that preserve accuracy tend to incur prohibitive compression overhead (Kim et al., 2025a). For instance, KVzip (Kim et al., 2025a) relies on the reconstruction process that doubles computation during prefill (see Figure 1b), limiting its practicality for latency-sensitive deployment.

In this paper, we introduce **Fast KVzip**, a fast and accurate KV cache eviction method that addresses the inefficiencies of prior reconstruction-based approaches. Our key insight is that the future utilization of KV pairs is largely an *intrinsic property* that can be directly decoded from the input hidden states, without reconstructing the entire context as in KVzip. Leveraging this insight, we seamlessly integrate a lightweight gating mechanism into the Transformer’s forward pass. This allows the model to directly assess the importance of KV pairs during both the prefill and decoding stages, incurring negligible memory and computational overhead while preserving high compression fidelity.

To identify an effective gating mechanism, we systematically explore the design space. Our empirical analysis reveals that gates perform most effectively when operating

on input hidden states rather than intermediate attention features. Furthermore, we find that an architecture inspired by attention sinks (Xiao et al., 2024) yields the strongest performance. Accordingly, we introduce a layer-wise gate training procedure that distills KV importance scores pre-computed from a compute-intensive reconstruction process. To ensure efficiency and avoid degrading the original model’s capabilities, we keep the LLM weights frozen and train only the gating parameters, requiring less than one H100 GPU hour for 14B-scale models.

As highlighted in Figure 1, Fast KVzip matches the compression performance of KVzip while maintaining model accuracy at KV budget ratios as low as 25% on the RULER-4K benchmark (Hsieh et al., 2024), and significantly outperforms existing baselines on the KVPress benchmark (as of December 2025) (Devoto et al., 2025). Moreover, Figure 1b shows that Fast KVzip substantially reduces peak memory usage and prefill time compared to both KVzip and the no-compression baseline.

We further demonstrate the effectiveness and generality of our method across a diverse set of state-of-the-art open-source LLMs, including Qwen2.5-7B/14B-1M (Yang et al., 2025b), Qwen3-8B/14B, Qwen3-8B-FP8 (Yang et al., 2025a), and Gemma3-12B (Team et al., 2025). Our approach is compatible with quantized weights and sliding-window hybrid attention mechanisms. We evaluate on prefill-intensive long-context benchmarks such as SCBench (Li et al., 2025) and MRCR (OpenAI, 2025) with context lengths of up to 170K tokens, as well as decoding-intensive reasoning tasks including AIME24 (AIME, 2025) and MATH (Hendrycks et al., 2021). Across all models and benchmarks, Fast KVzip achieves near-lossless performance while reducing the KV cache to 30–40% of its original size, underscoring the robustness and practicality of the proposed gating-based eviction strategy.

2. Related Work

KV Cache Compression. To mitigate the scalability bottleneck of KV caching in autoregressive Transformers (Radford et al., 2019), efficient attention mechanisms have been proposed, including sliding-window attention (Child et al., 2019), linear attention (Katharopoulos et al., 2020), latent attention (Liu et al., 2024a), and grouped-query attention (Ainslie et al., 2023). While effective, these methods typically require training from scratch or substantial fine-tuning to accommodate architectural changes.

Recent work observes inherent sparsity in LLM attention patterns (Liu et al., 2023). Methods such as H₂O (Zhang et al., 2023), Finch (Corallo & Papotti, 2024), and SnapKV (Li et al., 2024) exploit sparsity observed during prefill or decoding to evict under-utilized KV pairs. However, sparsity

Table 1. Gating strategies for Transformer attention. The table below indicates operation execution per layer for a given input. Mixture-of-Depths (MoD) uses gating to skip attention computation, while Mixture-of-Block-Attention (MoBA) applies gating to select blocks for sparse attention with stored KV features. Our method bridges these approaches by conducting an attention operation for every input token while improving computational and memory efficiency through gated KV eviction.

Method	Attention	Store KV feature
MoD (Raposo et al., 2024)	Conditional	Conditional
MoBA (Lu et al., 2025)	Always	Always
Fast KVzip (Ours)	Always	Conditional

patterns inferred at inference time often overfit the current input and fail to generalize across future queries. To address this limitation, query-agnostic approaches, *e.g.*, KVzip (Kim et al., 2025a), EpiCache (Kim et al., 2025b), and Expected Attention (Devoto et al., 2025), estimate KV importance independently of specific queries and demonstrate improved robustness in multi-query settings. Orthogonal strategies, such as DuoAttention (Xiao et al., 2025), RLKV (Du et al., 2025), and FastGen (Ge et al., 2024), predefine attention patterns (*e.g.*, sliding-window or A-shaped) using calibration datasets prior to deployment. Beyond token-level eviction, complementary compression techniques explore alternative dimensions, including feature-channel compression (Xu et al., 2025) and KV quantization (Liu et al., 2024b).

Gating in Transformers. To address the scaling issues of Transformers (Vaswani et al., 2017), researchers have introduced gating mechanisms that enable sparse and selective computation, including MoE (Shazeer et al., 2017), MoD (Raposo et al., 2024), and MoBA (Lu et al., 2025). These approaches demonstrate that pretraining with gating can effectively exploit sparsity in MLP and attention modules, improving computational efficiency. Table 1 provides a conceptual comparison of these approaches with our method.

In the landscape of KV cache compression, Locret (Huang et al., 2024) employs a gating module to predict KV importance scores, while DMS (Łańcucki et al., 2025) trains an eviction policy via logit distillation. Concurrent to our work, TrimKV (Bui et al., 2025) introduces a retention-gated mechanism to learn time-decaying importance scores. Despite their promise, Locret is confined to prefill-stage compression, and TrimKV suffers from performance degradation in prefill-intensive scenarios, specializing in narrow domains of mathematical reasoning. In our work, we provide a comprehensive interpretation of the gating mechanism, demonstrating that well-designed gating architectures and optimization strategies achieve near-lossless compression across both prefill and decoding stages. We provide a more detailed comparison to prior methods in Section C. We also note a concurrent work, KVzap (Jegou & Jeblick, 2026),

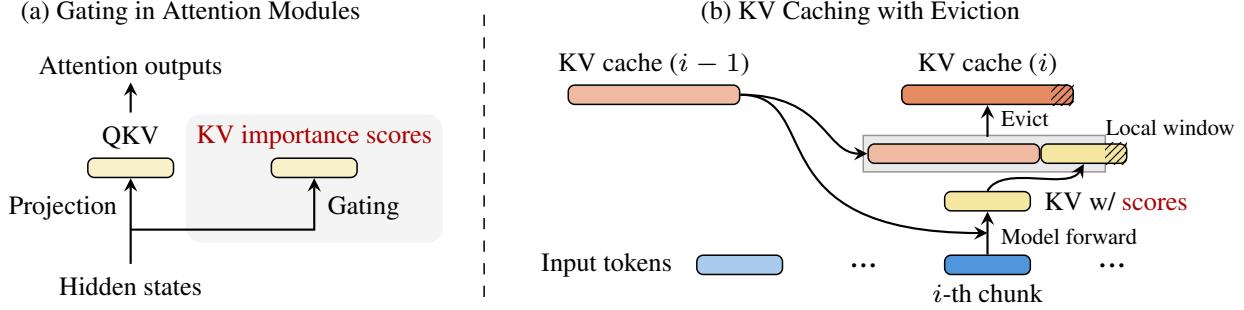


Figure 2. Computational flow. (a) Illustration of gating at each attention layer. (b) During each forward pass, we jointly calculate importance scores and evict low-importance KV pairs. For decoding, we maintain a small buffer to cache recent hidden states, performing gating and eviction in parallel once the buffer is full. This parallelized computation reduces latency overhead. We utilize a chunk size of 16K for prefill and a buffer size of 128 tokens for decoding; at each eviction step, we retain the KV pairs of the most recent tokens, keeping 4K tokens for prefill and 128 tokens for decoding.

which explores a similar gating optimization objective but differs in specific architecture and training recipe. Together with our work, these parallel efforts underscore the promise of gating-based approaches for KV cache compression.

3. Method

3.1. Gating Mechanism

Formulation. We design a gating module that processes the hidden features of each attention layer and outputs importance scores for the KV features in that layer. Specifically, the gate at layer l is defined as

$$g_l : \mathbb{R}^D \rightarrow [0, 1]^H,$$

where D denotes the hidden feature dimension and H is the number of KV heads in the layer. We introduce an additional gating branch into the model’s forward computation. This branch operates independently and does not affect the original attention outputs (see Figure 2a).

Computational Flow. Figure 2b illustrates the overall computational flow of our method. We adopt a chunked prefill to reduce the peak memory usage when processing long contexts (Agrawal et al., 2024). For each input chunk, we first compute the KV features and their corresponding importance scores using the gating mechanism. Based on these scores, we evict low-importance KV features from the cache, maintaining a compressed KV cache either at a fixed retention ratio or within a predefined memory budget. Our approach is compatible with both non-uniform and uniform KV cache structures (Feng et al., 2024; Li et al., 2024). Empirically, we observe that retaining KV features from recent tokens improves performance (see Section B.3). By maintaining a compressed KV cache throughout chunked prefill, our method reduces the peak KV cache memory compared to naive chunked prefill, as shown in Figure 1b.

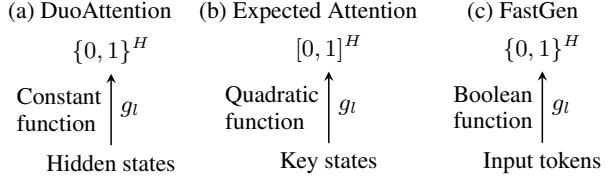


Figure 3. Gating formulation of existing KV compression methods at each attention layer. Note that we compare the core components of baseline methods, as each method jointly uses other techniques. For instance, FastGen combines its approach with cumulative attention scores (Zhang et al., 2023).

During decoding, we apply the same gated KV eviction strategy. However, computing gates at every decoding step introduces substantial overhead due to frequent function calls and increased I/O, resulting in a 30–60% latency increase depending on the context length. To mitigate this issue, we maintain a small buffer of token length 128 that caches recent hidden states and periodically computes gating in parallel. After updating the importance scores, we perform KV compression accordingly. This strategy reduces the gating overhead to approximately 1% of the model’s forward-pass latency on average.

3.2. Interpretation of Gating Formulation

Before detailing our method, we establish a conceptual basis for why gating mechanisms are effective for KV cache compression. Notably, we observe that several prominent compression methods can be reformulated as constrained instances of a generalized gating framework (Figure 3).

For instance, DuoAttention (Xiao et al., 2025) optimizes a head-level indicator function to designate specific heads for local-window attention. This represents a constrained gating formulation where the gating function is a constant Boolean mapping, independent of inputs. Expected Attention (De-

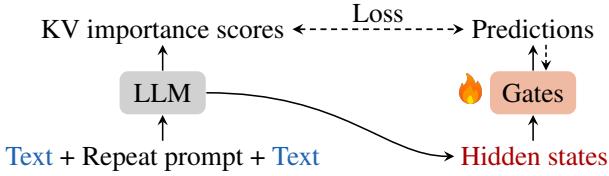


Figure 4. Gate training process. We compute the importance scores of KV pairs for a given text input following KVzip, which calculates the maximum attention score each KV pair receives during the context reconstruction (Kim et al., 2025a). We train gating modules to predict these scores from the input hidden states at each attention layer. Please refer to Section A for further details and the training hyperparameters.

voto et al., 2025) approximates future attention scores by applying a quadratic model of key states with mean query-state statistics. Furthermore, FastGen (Ge et al., 2024) identifies heads that attend to specific token clusters (e.g., punctuation or special tokens), which can be interpreted as a head-wise gating function on input tokens.

Our gating approach generalizes these prior methods by allowing any differentiable gating function without imposing empirical and heuristic constraints on the optimization. This interpretation enables a principled approach to KV cache compression, transforming the problem from heuristic design to data-driven optimization, effectively utilizing the model to discover the optimal compression strategy.

3.3. Gate Training

Consideration. We impose three requirements on our optimization: (1) computational efficiency, (2) robustness to task-specific overfitting, and (3) preservation of the base LLM’s performance. End-to-end optimization via backpropagation, for example, incurs substantial memory overhead, limiting the usable context length and increasing training time. Furthermore, training gating mechanisms on data from a specific domain or task, such as mathematical reasoning (Bui et al., 2025) or instruction-based question answering (Huang et al., 2024), may yield domain-specific gains but incur degraded performance on other domains or tasks.

Training Procedure. With these considerations, we propose a gate training procedure illustrated in Figure 4. We optimize the gates at each layer using target scores derived from KVzip’s reconstruction process (Kim et al., 2025a), while freezing LLM weights to avoid altering the model’s inherent capabilities. This design is computationally efficient, as reconstruction can be performed using parallelized forward passes. Given the target scores, we optimize gates independently at each layer using stochastic gradient descent with a binary cross-entropy loss (Wali, 2022). The resulting optimization is embarrassingly parallel across both samples and layers, enabling efficient GPU execution.

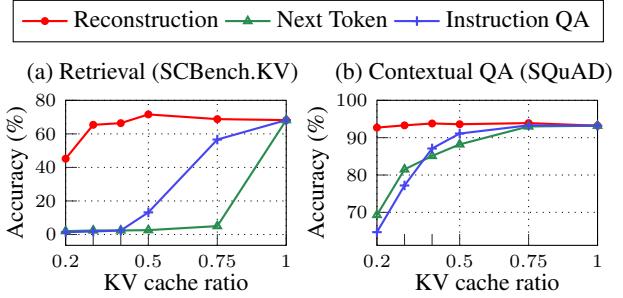


Figure 5. Effects of target scores for gate training, derived from reconstruction, next-token prediction, and instruction-based QA tasks. We evaluate Qwen2.5-7B-1M on a synthetic key retrieval task from SCBench (Li et al., 2025) and the SQuAD QA task (Rajpurkar et al., 2016).

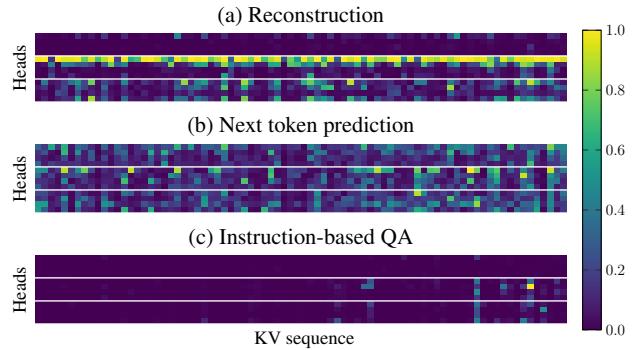


Figure 6. Maximum attention scores received by KV features across specific tasks. We visualize layers 7, 14, and 21 of Qwen2.5-7B-1M, with other layers exhibiting similar patterns, using the gate training data from Figure 5.

Reconstruction-based target signals have been shown to generalize well to downstream tasks (Devlin et al., 2019; Kim et al., 2025a), thereby mitigating task-overfitting issues. While KVzip introduces additional compression latency at inference time due to reconstruction, we perform this step only during a pre-deployment phase to train the gates. Consequently, inference remains fast.

Training Data. We train the gating mechanism using the FineWeb-Edu pretraining corpus (Penedo et al., 2024), which has no overlap with downstream evaluation datasets. We randomly sample sequences ranging from 10K to 30K tokens, yielding a total of 1M training tokens. Given that FineWeb-Edu contains over 1T tokens, our training set corresponds to approximately a 10^{-6} fraction of the full corpus.

Alternatives. In Figure 5, we compare the reconstruction-based target KV scores with scores derived from next-token prediction (Bui et al., 2025) and instruction-based question answering (Huang et al., 2024). We identically compute the maximum attention score received by each KV feature under

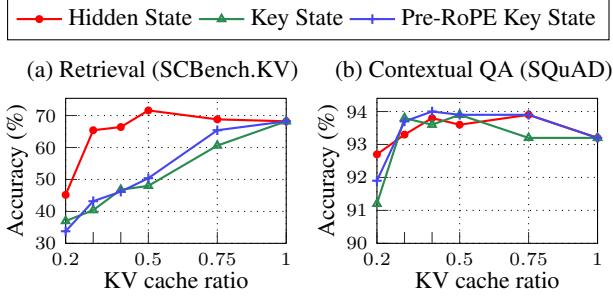


Figure 7. Effects of gate inputs, including hidden states, key states, and pre-RoPE key states, with the same settings as Figure 5.

each task. For instruction-based QA, we follow Huang et al. (2024) and use the LongAlpaca dataset (Chen et al., 2023). Figure 5 demonstrates that reconstruction-based targets generalize significantly better, while alternatives lead to inferior performance.

Figure 6 reveals that reconstruction-based attention patterns are uniform and structured across KV sequences. In contrast, next-token prediction generates denser patterns, reflecting more intensive feature contextualization. We hypothesize that the reconstruction task is effective because many downstream applications rely primarily on these high-level contextualized features (Kim et al., 2025a). To maintain necessary local dependencies, our method also preserves KV pairs within a sliding window (Figure 2). Conversely, QA tasks exhibit sparse attention as they focus only on query-specific information. Consequently, gates trained solely on QA signals generalize poorly, failing to retain the sufficient information required for broader downstream tasks.

In Figure 7, we study the choice of gate inputs by comparing hidden states, key states, and pre-RoPE key states (Su et al., 2024). Hidden states consistently outperform the alternatives, indicating that they contain richer information for predicting the future utility of KV features. We also observe that incorporating position-encoded features degrades performance. This motivates decoupling positional information from the gating mechanism and instead leveraging recency through a simple local-window strategy (Figure 2).

3.4. Architecture

Consideration. The primary objective in designing the gate architecture is to minimize computational and memory overhead while accurately modeling attention behavior to predict the future utilization of each KV feature.

Low-Rank Sink Attention. We introduce a lightweight, attention-inspired gating mechanism, illustrated in Figure 8. Given a hidden state $\mathbf{h} \in \mathbb{R}^D$, we apply linear projections followed by weighted normalization (Yang et al., 2025a)

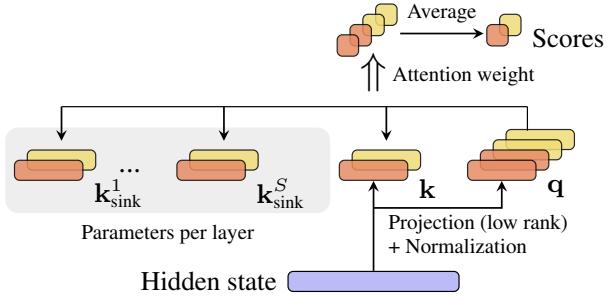


Figure 8. Gate architecture. We propose a sink-attention architecture using a fixed set of keys per layer. Given a hidden state $\mathbf{h} \in \mathbb{R}^D$, we compute the importance scores of KV as $\mathbf{s} \in [0, 1]^H$.

to obtain $\mathbf{k} \in \mathbb{R}^{H \times D'}$ and $\mathbf{q} \in \mathbb{R}^{G \times H \times D'}$, where H is the number of original KV heads and G denotes the original grouped-query size (Ainslie et al., 2023). Both representations are low-dimensional relative to the original KV features. The gate produces an importance score $\mathbf{s} \in \mathbb{R}^H$ via a sink-attention mechanism as

$$\mathbf{s} = \frac{1}{G} \sum_{j=1}^G \frac{\exp(\mathbf{q}_j^\top \mathbf{k})}{\sum_{r=1}^S \exp(\mathbf{q}_j^\top \mathbf{k}_{\text{sink}}^r) + b_j}.$$

Learnable sinks

Here, $\mathbf{k}_{\text{sink}}^r \in \mathbb{R}^{H \times D'}$ are layer-specific learnable parameters that model sink keys, which play a critical role in standard attention mechanisms (Xiao et al., 2024). We empirically observe that these features enable the gate to effectively estimate the importance of KV features for future attention. A scalar $b_j \geq 0$ serves as a bias controlling head-level importance. These parameters are trained jointly with the projection and normalization weights.

Unlike DeepSeek-V3.2, which employs a low-rank KV cache within its Lightning Indexer to enable sparse attention (Liu et al., 2025), our method uses a fixed set of low-rank key features per layer and introduces no additional KV caching. This design simplifies implementation and enables straightforward integration into existing architectures.

Note that we use an identical set of hyperparameters across all models, setting $S = 16$ and $D' = 16$, which is 1/8 of the Qwen3 head dimension (Yang et al., 2025a), which suggests that our method requires neither exhaustive hyperparameter search nor manual fine-tuning. We provide a sensitivity analysis of hyperparameters in Section B.

Alternatives. We study alternative gate architectures, including a linear model, a two-layer MLP with SwiGLU activation (Shazeer, 2020), and a variant of our architecture without a learnable denominator, *i.e.*, setting $\sum_{r=1}^S \exp(\mathbf{q}_j^\top \mathbf{k}_{\text{sink}}^r) + b_j = 1$. Except for the linear model,

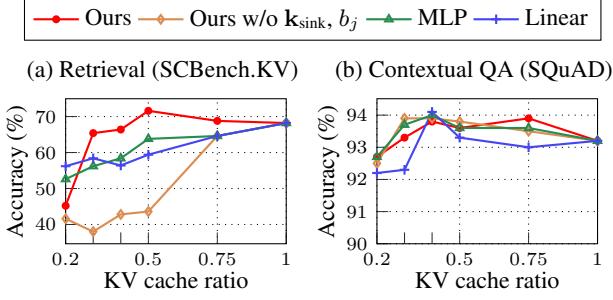


Figure 9. Performance of gate architectures. We compare our sink-attention, a variant without learnable denominator terms, a two-layer MLP, and a linear model, using Qwen2.5-7B-1M.

Table 2. Training efficiency. Gate training time and gate sizes across diverse model scales. The 30B-A3B model has a smaller hidden state size than the other models, resulting in smaller gates.

Model	Training Time (H100 hours)	Storage (GB)
Qwen3-8B	0.59	0.18
Qwen3-30B-A3B	0.70	0.11
Qwen3-14B	0.83	0.30

we match the parameter counts across all architectures. Figure 9 shows that while alternative architectures are competitive on contextual QA tasks, they exhibit degraded performance on retrieval tasks. These results highlight the effectiveness of our sink-attention gate architecture.

3.5. Efficiency

Gate Training. Table 2 shows the training time and memory requirements of our gate training process, demonstrating its efficiency across diverse model scales.

Inference. Figure 10 shows the memory and latency efficiency improvements of Fast KVzip during both prefill and decoding. The results demonstrate that Fast KVzip reduces latency and memory usage compared to the base model, while significantly improving prefilling efficiency compared to KVzip by removing the runtime KV cache compression overhead. For all experiments, we apply a non-uniform KV cache eviction structure following Kim et al. (2025a).

4. Experiment

We evaluate the KV cache compression performance of Fast KVzip under both prefill-intensive and decoding-intensive scenarios. We train gates once per model using FineWeb-Edu (Penedo et al., 2024) and evaluate on all downstream tasks (see Section 3.3), demonstrating that the learned gates are task-generalizable. All gate training and model evaluations are conducted on a single NVIDIA H100 80GB GPU using PyTorch and FlashAttention-2 (Dao, 2024).

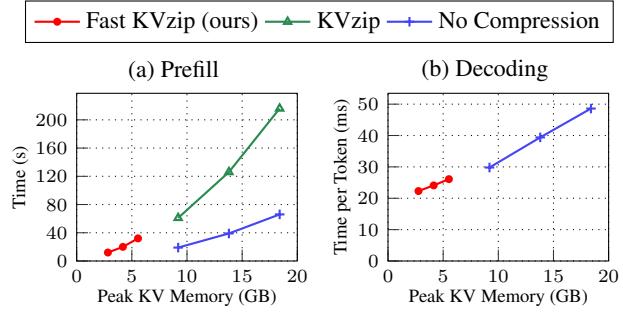


Figure 10. Inference efficiency. Prefill and decoding efficiency of the Qwen2.5-7B-1M model using a 30% KV budget ratio with PyTorch and FlashAttention-2 (Dao, 2024) on a single H100 GPU. Points in the plot correspond to context lengths of 160K, 240K, and 320K. Note that KVzip provides a decoding speed similar to Fast KVzip, which is omitted from Figure (b) for clarity.

4.1. Setup

Datasets. For prefill-intensive scenarios, we evaluate all methods on SCBench (Li et al., 2025), a multi-query benchmark comprising nine tasks, including synthetic key retrieval tasks from RULER (Hsieh et al., 2024), long-context question answering from LongBench (Bai et al., 2024), and code comprehension tasks. We further include OpenAI MRCR (OpenAI, 2025), a multi-round conversational retrieval benchmark, and SQuAD (Rajpurkar et al., 2016), a multi-query QA dataset. The maximum context length reaches up to 170K tokens using the Qwen3 tokenizer.

For decoding-intensive tasks, we evaluate on mathematical reasoning benchmarks, AIME24 (AIME, 2025) and MATH (Hendrycks et al., 2021). Following the R-KV baseline (Cai et al., 2025), we set the maximum decoding length to 32K tokens for AIME24 and 16K tokens for MATH.

Models. We conduct experiments on a diverse set of state-of-the-art open-source large language models, including Qwen2.5-7B/14B-1M (Yang et al., 2025b), Qwen3-8B/14B, Qwen3-8B-FP8 (Yang et al., 2025a), and Gemma3-12B (Team et al., 2025). We evaluate each model using its native precision. Notably, Gemma3-12B employs a hybrid sliding-window attention mechanism, while Qwen3-8B-FP8 adopts dynamic FP8 quantization. For the hybrid model, we only compress the global attention KV cache, which dominates memory usage in long-context evaluation.

Baselines. We compare Fast KVzip against several state-of-the-art KV cache compression methods, including KVzip (Kim et al., 2025a), SnapKV (Li et al., 2024), Expected Attention (Devoto et al., 2025) with Ada-KV (Feng et al., 2024), DuoAttention (Xiao et al., 2025) with KVzip score (Kim et al., 2025a), R-KV (Cai et al., 2025), and TrimKV (Bui et al., 2025). For prefill-intensive tasks, we evaluate

Fast KVzip

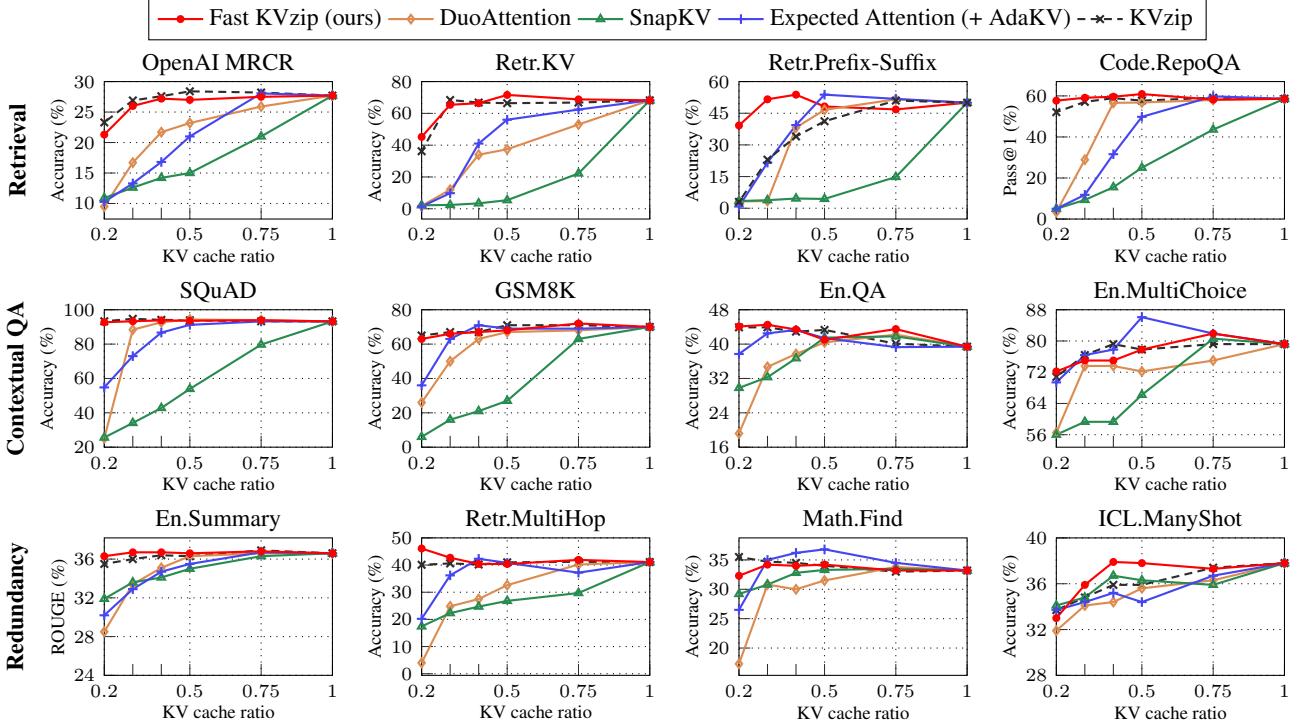


Figure 11. **Prefill-intensive benchmark results** with Qwen2.5-7B-1M under varying KV cache budget ratios from 0.2 to 1.0. Tasks are organized into three categories: (1) retrieval-intensive, (2) contextual understanding, and (3) high context redundancy.

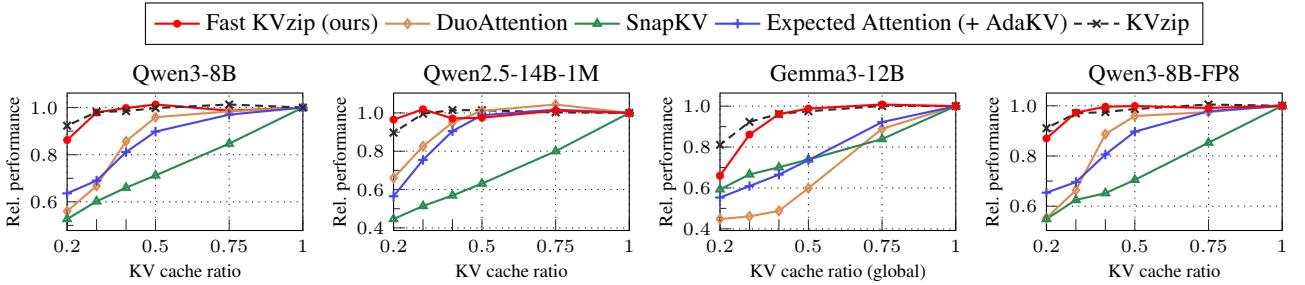


Figure 12. **Performance across different models**, averaged over 12 benchmark datasets. For each dataset, results are normalized with respect to full-cache performance prior to averaging.

methods under the query-agnostic setting, following the setting of KVzip (Kim et al., 2025a).

For decoding-intensive tasks, we adopt the R-KV setting (Cai et al., 2025), which continuously compresses the KV cache during decoding, covering both the context and generated tokens, to match the fixed-size cache budget. Additionally, we evaluate the early stopping of thinking strategy introduced in Qwen3 (Yang et al., 2025a), where the model terminates the thinking process upon reaching a predefined budget and is forced to generate the final answer. Since some baselines are specifically designed for either prefill-intensive or decoding-intensive scenarios, we report comparisons only within their respective applicable settings.

4.2. Prefill-Intensive Tasks

Figure 11 shows the prefill KV cache compression performance across 12 datasets. The results demonstrate that Fast KVzip outperforms all baselines across all tasks while matching the performance of KVzip, which requires twice the prefilling computational cost for compression. The figure shows that Fast KVzip maintains full-cache performance at a 30-40% KV cache budget ratio. On some tasks, Fast KVzip exhibits performance improvements, which we attribute to the denoising effect of attention induced by KV cache compression (Ye et al., 2025).

Figure 12 presents the averaged relative performance across the 12 datasets for various LLMs, demonstrating that Fast KVzip maintains effectiveness across models with

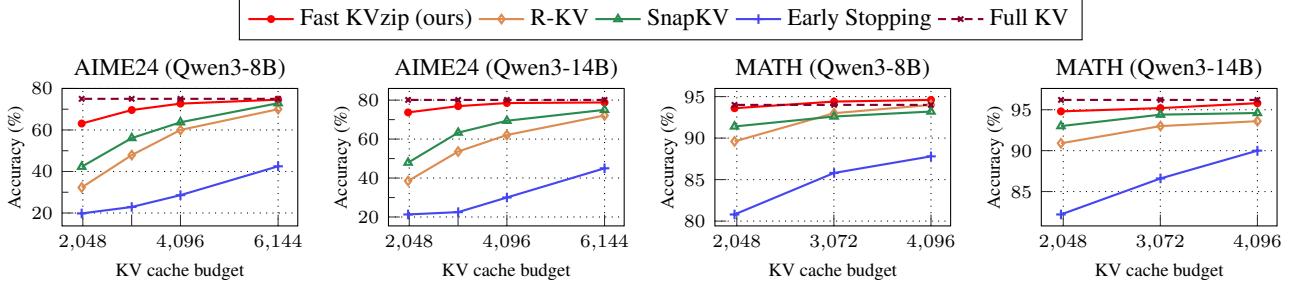


Figure 13. Decoding-intensive benchmark results. For AIME24, we report the average score over 16 random seeds.

a different pre-/post-training (Qwen3-8B), a larger model scale (Qwen2.5-14B-1M), a different attention mechanism (Gemma3-12B), and a quantization (Qwen3-8B-FP8).

4.3. Decoding-Intensive Tasks

Figure 13 shows the KV cache compression performance on mathematical reasoning tasks, including AIME24 and MATH. The results demonstrate that Fast KVzip achieves superior performance compared to decoding-scenario compression baselines on Qwen3-8B and Qwen3-14B models. Fast KVzip achieves near-lossless performance with a KV budget size of 4K, whereas the early-stopping-of-thinking method drastically reduces reasoning performance. These results indicate that the model requires a sufficient length of thinking to deduce correct answers, and that Fast KVzip indeed maintains the necessary KV features for general inference, including reasoning procedures.

5. Qualitative Analysis

We analyze the gating behavior of Qwen2.5-7B-1M using English book and code data from SCBench. Figure 14 illustrates the distribution of KV cache retention rates across attention heads under a total of 36% KV cache budget. We observe that half of the heads (primarily in early layers) exhibit highly sparse KV caches, while the remaining heads show uniformly distributed retention rates. This suggests that the gates prioritize high-level features from middle-to-late layers over low-level features from early layers.

We categorize attention heads by their retention rates (r) into three types: sparse ($r < 0.05$), medium ($0.05 \leq r < 0.9$), and dense ($0.9 \leq r$). Table 3 details the tokens most frequently evicted and retained by each head type. Notably, sparse heads predominantly retain tokens associated with punctuation or line-break tokens. In contrast, other heads tend to retain tokens that are frequent in the source data; these frequent tokens appear as both the most retained and most evicted in medium heads. We hypothesize that sparse heads capture high-level, summarized information (Razzhigaev et al., 2025), indicating an attention pattern distinct

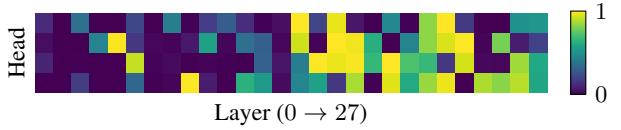


Figure 14. Retention Rates per Attention Head. We apply a 36% KV cache budget to Qwen2.5-7B-1M, which has 28 layers with 4 KV heads per layer. The left side corresponds to the early layers.

Table 3. Gate behavior analysis. The most frequently retained and evicted tokens for each head type. The \s represents a space.

Head Type	Most Retained (Top 1 → 10)	Most Evicted (Top 1 → 10)
Sparse	. \n\n\n , for \n\n\n all that	, the . and to of a I in was
Medium	, . and the of to I in \n a	the , and to of a . I in was
Dense	, the . and to of a I in was	the \s , a \t \s\s (1 of I

from local-window mechanisms (Xiao et al., 2024).

6. Conclusion

Summary. We propose Fast KVzip, a gating-based KV cache compression method that matches state-of-the-art performance while eliminating runtime compression overhead. Through principled analysis, we design an effective gating optimization and a lightweight sink-attention architecture. Extensive evaluations on both prefill- and decoding-intensive tasks show that Fast KVzip generalizes well across diverse downstream applications, including retrieval, code comprehension, and reasoning, reducing KV cache size to 30% with negligible performance loss.

Limitations and Future Work. Prior gating methods for attention modules, such as MoD (Raposo et al., 2024) and MoBA (Lu et al., 2025), typically require from-scratch training to fit predefined structures. An important future direction is to study how applying our method during pretraining can induce hardware-efficient structured computation patterns. Moreover, exploring a unified framework that integrates these gating approaches is a promising avenue. In particular, the gating mechanism could be extended to support multi-choice decisions, including skipping computation, discarding KV features, or storing them for selective retrieval.

References

Agrawal, A., Kedia, N., Panwar, A., Mohan, J., Kwatra, N., Gulavani, B., Tumanov, A., and Ramjee, R. Taming throughput-latency tradeoff in llm inference with sarathi-serve. *18th USENIX Symposium on Operating Systems Design and Implementation (OSDI 24)*, 2024.

AIME. Aime problems and solutions, 2025. URL https://artofproblemsolving.com/wiki/index.php/AIME_Problems_and_Solutions.

Ainslie, J., Lee-Thorp, J., De Jong, M., Zemlyanskiy, Y., Lebrón, F., and Sanghai, S. Gqa: Training generalized multi-query transformer models from multi-head checkpoints. *EMNLP*, 2023.

Bai, Y., Lv, X., Zhang, J., Lyu, H., Tang, J., et al. Long-bench: A bilingual, multitask benchmark for long context understanding. *ACL*, 2024.

Bui, N., Sharma, S., Lamba, S., Mishra, S., and Ying, R. Cache what lasts: Token retention for memory-bounded kv cache in llms. *arXiv:2512.03324*, 2025.

Cai, Z., Xiao, W., Sun, H., Luo, C., Zhang, Y., et al. R-kv: Redundancy-aware kv cache compression for training-free reasoning models acceleration. *Advances in Neural Information Processing Systems*, 2025.

Chen, Y., Qian, S., Tang, H., Lai, X., Liu, Z., Han, S., and Jia, J. Longlora: Efficient fine-tuning of long-context large language models. *arXiv:2309.12307*, 2023.

Child, R., Gray, S., Radford, A., and Sutskever, I. Generating long sequences with sparse transformers. *arXiv:1904.10509*, 2019.

Comanici, G., Bieber, E., Schaeckermann, M., Pasupat, I., et al. Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities. *arXiv:2507.06261*, 2025.

Corallo, G. and Papotti, P. Finch: Prompt-guided key-value cache compression for large language models. *Transactions of the Association for Computational Linguistics*, 2024.

Dai, Z., Yang, Z., Yang, Y., Carbonell, J. G., Le, Q., and Salakhutdinov, R. Transformer-xl: Attentive language models beyond a fixed-length context. *Proceedings of the 57th annual meeting of the association for computational linguistics*, 2019.

Dao, T. Flashattention-2: Faster attention with better parallelism and work partitioning. *ICLR*, 2024.

Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. Bert: Pre-training of deep bidirectional transformers for language understanding. *NAACL*, 2019.

Devoto, A., Jeblick, M., and Jégou, S. Expected attention: Kv cache compression by estimating attention from future queries distribution. *arXiv:2510.00636*, 2025. URL <https://arxiv.org/abs/2510.00636>.

Du, W., Jiang, L., Tao, K., Liu, X., and Wang, H. Which heads matter for reasoning? rl-guided kv cache compression. *arXiv preprint arXiv:2510.08525*, 2025.

Feng, Y., Lv, J., Cao, Y., Xie, X., and Zhou, S. K. Adakv: Optimizing kv cache eviction by adaptive budget allocation for efficient llm inference. *arXiv:2407.11550*, 2024.

Ge, S., Zhang, Y., Liu, L., Zhang, M., Han, J., and Gao, J. Model tells you what to discard: Adaptive kv cache compression for llms. *ICLR*, 2024.

Hendrycks, D., Burns, C., Kadavath, S., Arora, A., Basart, S., Tang, E., Song, D., and Steinhardt, J. Measuring mathematical problem solving with the math dataset. *arXiv:2103.03874*, 2021.

Hsieh, C.-P., Sun, S., Kriman, S., Acharya, S., Rekesh, D., Jia, F., Zhang, Y., and Ginsburg, B. Ruler: What's the real context size of your long-context language models? *COLM*, 2024.

Huang, Y., Yuan, B., Han, X., Xiao, C., and Liu, Z. Locret: Enhancing eviction in long-context llm inference with trained retaining heads on consumer-grade devices. *arXiv:2410.01805*, 2024.

Jaech, A., Kalai, A., Lerer, A., Richardson, A., El-Kishky, A., et al. Openai o1 system card. *arXiv:2412.16720*, 2024.

Jegou, S. and Jeblick, M. Kvzap: Fast, adaptive, and faithful kv cache pruning. *arXiv:2601.07891*, 2026.

Katharopoulos, A., Vyas, A., Pappas, N., and Fleuret, F. Transformers are rnns: Fast autoregressive transformers with linear attention. *ICML*, 2020.

Kim, J.-H., Kim, J., Kwon, S., Lee, J. W., Yun, S., and Song, H. O. Kvzip: Query-agnostic kv cache compression with context reconstruction. *Advances in Neural Information Processing Systems*, 2025a.

Kim, M., Kundu, A., Kim, H.-B., Dixit, R., and Cho, M. Epi-cache: Episodic kv cache management for long conversational question answering. *arXiv:2509.17396*, 2025b.

Kwon, W., Li, Z., Zhuang, S., Sheng, Y., Zheng, L., Yu, C. H., Gonzalez, J., Zhang, H., and Stoica, I. Efficient memory management for large language model serving with pagedattention. *Proceedings of the 29th Symposium on Operating Systems Principles*, 2023.

Łańcucki, A., Staniszewski, K., Nawrot, P., and Ponti, E. M. Inference-time hyper-scaling with kv cache compression. *Advances in Neural Information Processing Systems*, 2025.

Li, Y., Huang, Y., Yang, B., Venkitesh, B., Locatelli, A., Ye, H., Cai, T., Lewis, P., and Chen, D. Snapkv: Llm knows what you are looking for before generation. *Advances in Neural Information Processing Systems*, 2024.

Li, Y., Jiang, H., Wu, Q., Luo, X., Ahn, S., Zhang, C., Abdi, A. H., Li, D., Gao, J., Yang, Y., et al. Scbench: A kv cache-centric analysis of long-context methods. *ICLR*, 2025.

Liu, A., Feng, B., Wang, B., Wang, B., Liu, B., et al. Deepseek-v2: A strong, economical, and efficient mixture-of-experts language model. *arXiv:2405.04434*, 2024a.

Liu, A., Mei, A., Lin, B., Xue, B., Wang, B., Xu, B., Wu, B., Zhang, B., Lin, C., Dong, C., et al. Deepseek-v3. 2: Pushing the frontier of open large language models. *arXiv:2512.02556*, 2025.

Liu, Z., Wang, J., Dao, T., Zhou, T., Yuan, B., Song, Z., et al. Deja vu: Contextual sparsity for efficient llms at inference time. *ICML*, 2023.

Liu, Z., Yuan, J., Jin, H., Zhong, S., Xu, Z., Braverman, V., Chen, B., and Hu, X. Kivi: A tuning-free asymmetric 2bit quantization for kv cache. *ICML*, 2024b.

Lu, E., Jiang, Z., Liu, J., Du, Y., Jiang, T., Hong, C., et al. Moba: Mixture of block attention for long-context llms. *arXiv:2502.13189*, 2025.

OpenAI. Openai mrcr: Long context multiple needle in a haystack benchmark, 2025. URL <https://huggingface.co/datasets/openai/mrcr>.

Penedo, G., Kydlíček, H., Lozhkov, A., Mitchell, M., Raffel, C. A., Von Werra, L., Wolf, T., et al. The fineweb datasets: Decanting the web for the finest text data at scale. *Advances in Neural Information Processing Systems*, 37: 30811–30849, 2024.

Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., Sutskever, I., et al. Language models are unsupervised multitask learners. *OpenAI blog*, 2019.

Rajpurkar, P., Zhang, J., Lopyrev, K., and Liang, P. Squad: 100,000+ questions for machine comprehension of text. *EMNLP*, 2016.

Raposo, D., Ritter, S., Richards, B., Lillicrap, T., Humphreys, P. C., and Santoro, A. Mixture-of-depths: Dynamically allocating compute in transformer-based language models. *arXiv:2404.02258*, 2024.

Razzhigaev, A., Mikhalkhuk, M., Rahmatullaev, T., Goncharova, E., Druzhinina, P., Oseledets, I., and Kuznetsov, A. Llm-microscope: Uncovering the hidden role of punctuation in context memory of transformers. *arXiv:2502.15007*, 2025.

Shazeer, N. Glu variants improve transformer. *arXiv:2002.05202*, 2020.

Shazeer, N., Mirhoseini, A., Maziarz, K., Davis, A., Le, Q., Hinton, G., and Dean, J. Outrageously large neural networks: The sparsely-gated mixture-of-experts layer. *arXiv:1701.06538*, 2017.

Su, J., Ahmed, M., Lu, Y., Pan, S., Bo, W., and Liu, Y. Roformer: Enhanced transformer with rotary position embedding. *Neurocomputing*, 2024.

Team, G., Kamath, A., Ferret, J., Pathak, S., Vieillard, N., et al. Gemma 3 technical report. *arXiv:2503.19786*, 2025.

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., and Polosukhin, I. Attention is all you need. *Advances in Neural Information Processing Systems*, 2017.

Wali, R. Xtreme margin: A tunable loss function for binary classification problems. *arXiv:2211.00176*, 2022.

Xiao, G., Tian, Y., Chen, B., Han, S., and Lewis, M. Efficient streaming language models with attention sinks. *ICLR*, 2024.

Xiao, G., Tang, J., Zuo, J., Guo, J., Yang, S., Tang, H., Fu, Y., and Han, S. Duoattention: Efficient long-context llm inference with retrieval and streaming heads. *ICLR*, 2025.

Xu, Y., Jie, Z., Dong, H., Wang, L., Lu, X., Zhou, A., Saha, A., Xiong, C., and Sahoo, D. Think: Thinner key cache by query-driven pruning. *The Thirteenth International Conference on Learning Representations*, 2025.

Yang, A., Li, A., Yang, B., Zhang, B., Hui, B., et al. Qwen3 technical report. *arXiv:2505.09388*, 2025a.

Yang, A., Yu, B., Li, C., Liu, D., Huang, F., Huang, H., et al. Qwen2.5-1m technical report. *arXiv:2501.15383*, 2025b.

Ye, T., Dong, L., Xia, Y., Sun, Y., Zhu, Y., Huang, G., and Wei, F. Differential transformer. *ICLR*, 2025.

Zhang, Z., Sheng, Y., Zhou, T., Chen, T., Zheng, L., et al. H2o: Heavy-hitter oracle for efficient generative inference of large language models. *Advances in Neural Information Processing Systems*, 2023.

A. Experimental Details

In this section, we describe the training hyperparameters used in Section 3.3 and the gate architecture detailed in Section 3.4.

A.1. Dataset

As described in Section 3.3, we use the FineWeb-Edu dataset (Penedo et al., 2024) to train the gates. We randomly subsample sequences with context lengths ranging from 10K to 30K tokens, for a total of 500K tokens. By concatenating these samples, we construct an additional long-context dataset, yielding sequences of length 100K tokens and totaling another 500K tokens. In total, we use 1M tokens for gate training.

A.2. Training Hyperparameters

For efficiency, we first precompute tuples of hidden states and their corresponding KV importance scores using the dataset described above. We then optimize the gates in parallel using stochastic gradient descent with a binary cross-entropy loss and a learning rate of 0.2 across all experiments. Training is performed for 5K update steps with a batch size of 1K. Given the total of 1M training tokens, this corresponds to 5 epochs over the dataset.

Figure 15 illustrates the training loss trajectory and demonstrates stable convergence. We observe that earlier layers tend to exhibit lower loss, whereas the middle and later layers show higher loss values. We conjecture that higher layers perform more complex attention mechanisms, which are inherently more difficult to predict from a single hidden state.

Figure 20 presents a token-level comparison of predictions, illustrating that our trained gating mechanism captures token-level dynamics and avoids over-smoothing.

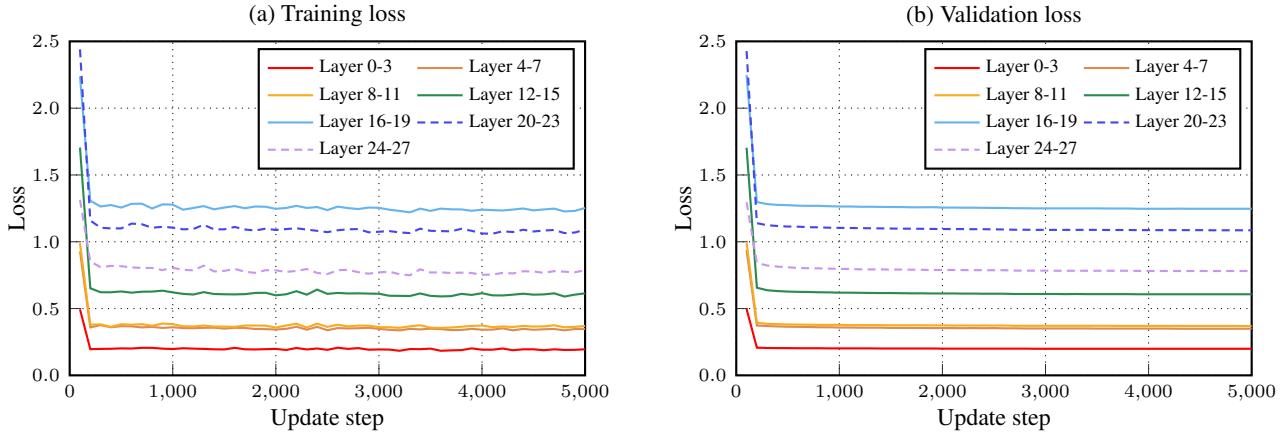


Figure 15. Gate training and validation loss trajectory per attention layer for Qwen2.5-7B-1M. We plot loss values averaged every four layers in the model.

A.3. Gate Configuration

As described in Section 3.4, each gate uses 16 $\mathbf{k}_{\text{sink}}^s$ vectors. The low-rank projection dimension D' is set to 16 for all experiments.

A.4. Inference Details

During prefill, we employ chunked prefill with a chunk size of 16K tokens. As illustrated in Figure 2, KV cache eviction is performed after each prefill chunk, while retaining KV features from the most recent 4K-token window to exploit recency bias in attention. For contexts shorter than 16K tokens, we retain the last 2% of tokens to accommodate arbitrary input lengths.

During decoding, KV eviction begins once the cache reaches the predefined budget. We use a hidden cache of size 128 and compute gating decisions in parallel. After updating importance scores, we evict an equivalent number of KV entries to maintain the cache budget. A local attention window of 128 tokens is preserved at each eviction step during decoding.

B. Hyperparameter Sensitivity Analysis

B.1. Training Data

Figure 16 shows the performance of gates trained with different sizes of FineWeb-Edu data, as described in Section A.1. The figure demonstrates that increasing the size of the training data leads to performance improvements, particularly on SCBench.KV, the synthetic key-retrieval task. To further evaluate the effect of long-context data, we construct the 500K-short dataset by removing the long-context samples (each with a context length of 100K) from the 1M-token dataset described in Section A.1. The results demonstrate the effectiveness of incorporating long-context data for gate training.

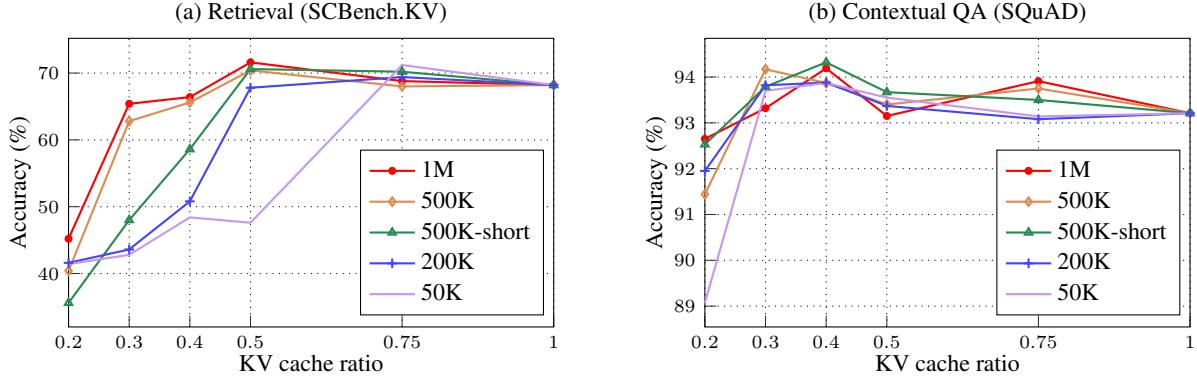


Figure 16. **Effect of Training Data Size.** We evaluate gates trained with different amounts of training data sampled from FineWeb-Edu using Qwen2.5-7B-1M. For the other settings, we uniformly reduce the context lengths of the data.

B.2. Gate Configuration

Figure 17 analyzes the effects of the number of sink keys S and the projection dimension D' in our sink-attention architecture (Section 3.4). The figure shows that performance is robust to the choice of projection dimension, while highlighting the importance of having a sufficient number of sink keys for effective KV cache compression.

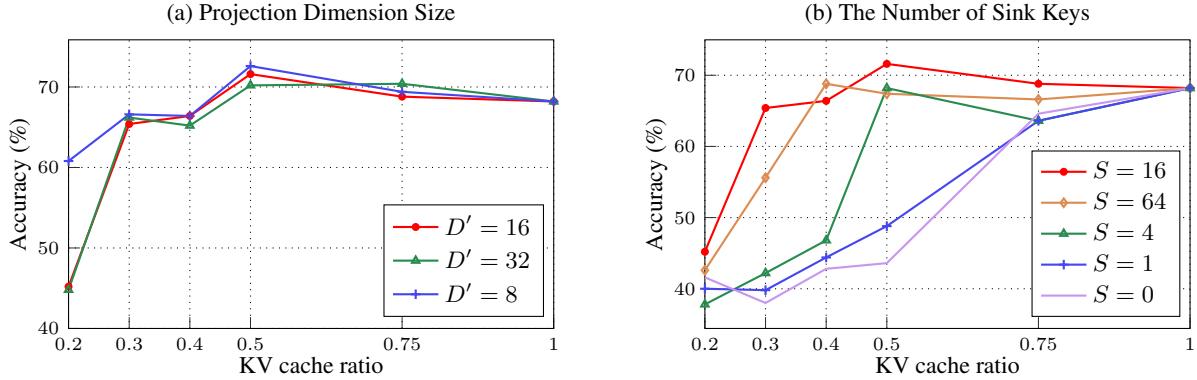


Figure 17. **Effect of Gate Configuration.** We evaluate gates with different numbers of sink keys S and projection dimensions D' using Qwen2.5-7B-1M on SCBench.KV.

B.3. Inference Configuration

Figure 18 shows the effect of the local window size in our inference algorithm described in Figure 2. The figure demonstrates that retaining a local window improves performance, while performance is comparable across local window sizes ranging from 1K to 8K tokens.

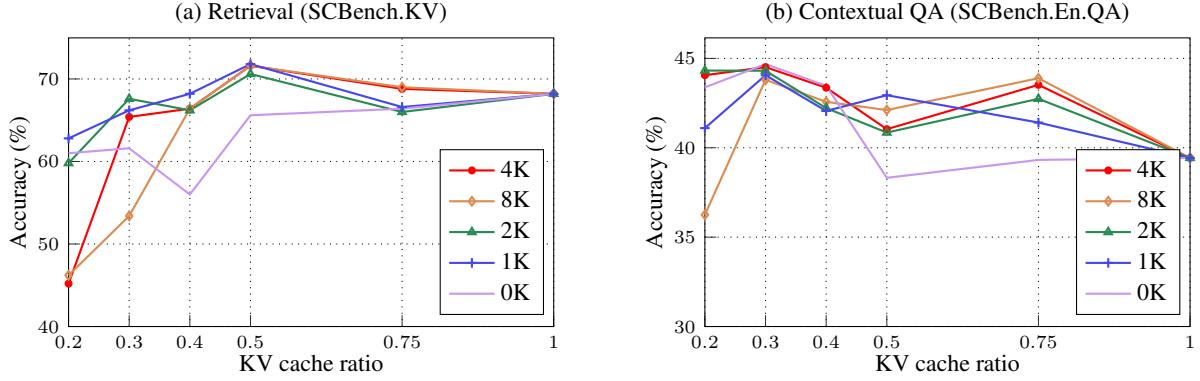


Figure 18. **Effect of Local Window Size.** We evaluate gates with a range of local window sizes, as described in Figure 2, using Qwen2.5-7B-1M. Note that our algorithm performs eviction that all configurations have an identical total KV cache size after eviction.

C. Extended Comparison with Prior Gating Approaches

In Table 4, we provide a comparison with prior gating-based approaches for KV cache compression, including Locret (Huang et al., 2024), DMS (Łańcucki et al., 2025), and TrimKV (Bui et al., 2025).

In Figure 19, we further evaluate our method on Qwen3-4B-Instruct-2507, for which an official TrimKV model for general language tasks is available. For other released models, TrimKV trains specialized gating models for mathematical reasoning using math-specific datasets, which prevents a fair comparison with our method, trained on general data in a task-agnostic manner. As shown in Figure 19, our Fast KVzip outperforms TrimKV in general task settings, particularly on retrieval tasks. These results demonstrate the effectiveness of our gating optimization and design, emphasizing their importance.

Table 4. **Comparison of gating approaches** for KV cache compression. While DMS and TrimKV trains separate gates for reasoning and instruction-following tasks using distinct datasets, we employ a single dataset for gate training across all models and evaluations.

Method	Training Data	Task	Target	Gate Architecture
Locret (Huang et al., 2024)	Instruction-based QA	Question-answering	Attention scores	MLP
DMS (Łańcucki et al., 2025)	Mathematics (or Programming)	Next-token logit distillation	End-to-end loss	Linear
TrimKV (Bui et al., 2025)	Mathematics (or Book)	Next-token prediction	End-to-end loss	MLP
Fast KVzip (Ours)	General text data	Reconstruction	Attention scores	Sink attention

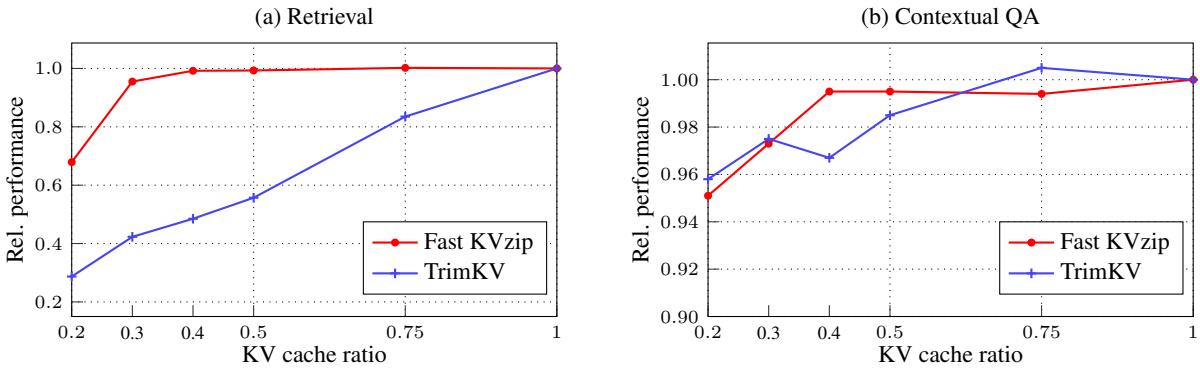


Figure 19. **Comparison to TrimKV** (Bui et al., 2025) with prefill-intensive tasks. We evaluate the official TrimKV model with Qwen3-4B-Instruct-2507 on retrieval and contextual QA tasks in Figure 11, and report the averaged scores. For each dataset, results are normalized with respect to full-cache performance prior to averaging. For Locret, the released models are outdated, and we replace performance comparisons with our analysis in Figure 5.

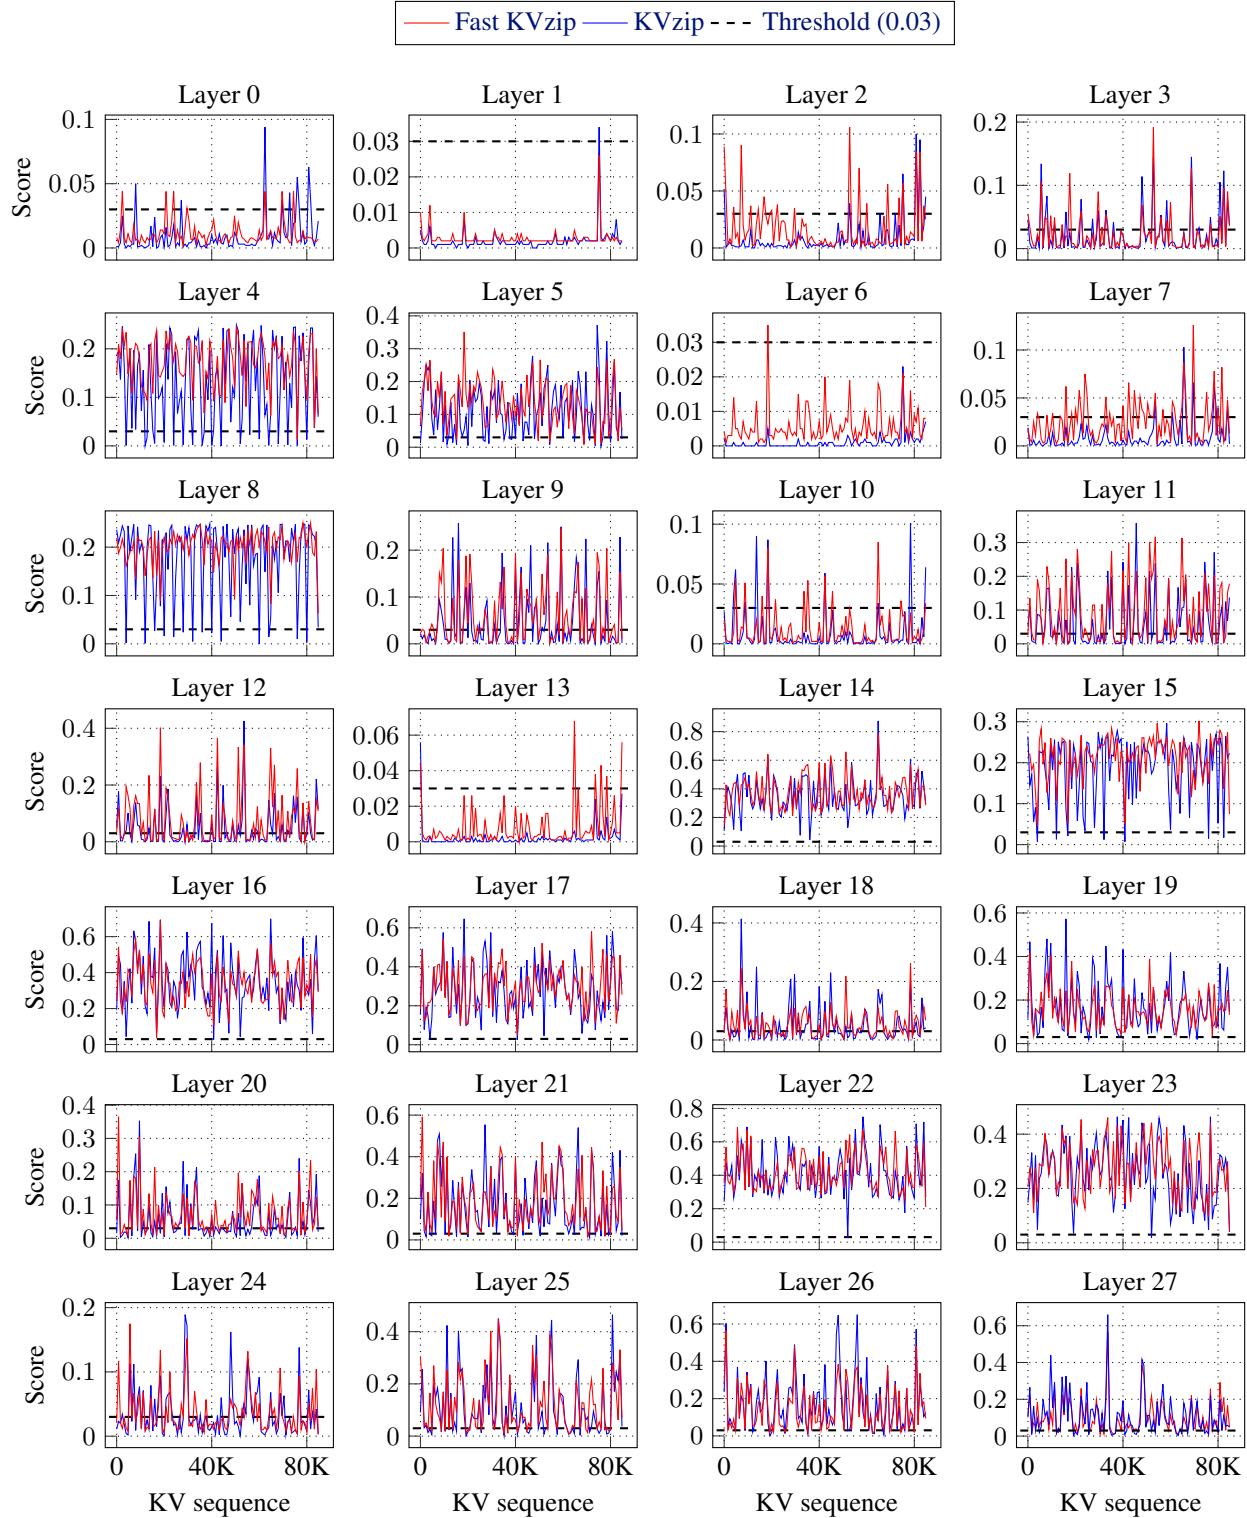


Figure 20. Prediction Comparison of Fast KVzip to KVzip. We use Qwen2.5-7B-1M with the SCBench.En.QA dataset, which were not used during training. For each layer, we plot the averaged score over heads. The dashed line represents the threshold value that results in a 36% KV cache budget ratio.