SeqFusion: Scalable Long-Context Reasoning through Parallel Fragment Fusion and Memory-Augmented Attention

Yanxuan Yu*
Columbia University
New York, NY
yy3523@columbia.edu

Dong Liu*
Yale University
New Haven, CT
dong.liu.dl2367@yale.edu

Abstract

Large Language Models (LLMs) often exhibit superior performance on short-context inference ($\leq$2k tokens) compared to long-context reasoning ($\geq$32k tokens), a phenomenon we term the *fragmentation gap*. This gap stems from training-inference mismatch and cumulative attention drift in long sequences. We propose **SeqFusion**, a novel framework that bridges fragmented short-context inference with unified long-context reasoning through consistency alignment and memory linking. SeqFusion achieves **2-3x speedup** in long-context processing while maintaining or improving accuracy through its innovative fragmented inference approach. Our method introduces fragment-to-long alignment loss and cross-fragment memory anchors, enabling models to leverage the accuracy benefits of short-context inference while maintaining global consistency. Extensive experiments on LongBench, BookSum, and Passkey Retrieval demonstrate that SeqFusion significantly reduces tail degradation (TDS improvement of 0.15-0.25) and increases fragment-long consistency (FLC improvement of 0.2-0.35) while achieving **40-60% memory reduction** and **2.5-3.5x throughput improvement** compared to traditional long-context approaches.

1 Introduction

Large Language Models (LLMs) have revolutionized natural language processing, achieving remarkable performance across a wide range of tasks. However, these models face significant challenges when processing long-context sequences, particularly those exceeding their training context window. While models excel at short-context inference ($\leq$2k tokens), their performance degrades substantially on long sequences ($\geq$32k tokens), exhibiting what we identify as the *fragmentation gap*. This gap manifests through memory degradation, attention drift, training-inference mismatch, and computational inefficiency, particularly problematic in real-world applications such as document analysis, code understanding, multi-turn conversations, and information retrieval requiring processing of large knowledge bases.

Traditional approaches including KV compression, sliding window attention, recurrent memory, and hierarchical processing address symptoms rather than the root cause of the fragmentation gap. Our key insight is that parallel processing of short-context fragments with strategic memory linking can achieve superior performance to unified long-context reasoning. This stems from three principles: models are optimized for short-context scenarios, short-context inference avoids attention drift, and fragment-level processing leverages model strengths while avoiding weaknesses. However,

*Equal contribution

39th Conference on Neural Information Processing Systems (NeurIPS 2025) Workshop: ER @ NeurIPS 2025 Poster.

independent fragment processing leads to loss of global consistency, inability to leverage cross-fragment information, and potential contradictions between outputs.

SeqFusion addresses these challenges through parallel fragment distribution across multiple GPUs with dynamic load balancing, memory-augmented attention with cross-fragment memory anchors for global consistency, distributed memory synchronization through AllReduce-based memory management, and CUDA-optimized processing with custom kernels achieving optimal performance. This paper makes five key contributions: scalable parallel processing enabling linear scaling across 2-8 GPUs with 90% efficiency at 4 GPUs, memory-augmented attention with distributed memory anchors maintaining global consistency, CUDA-optimized implementation achieving 1.35-1.51x performance improvement, comprehensive evaluation demonstrating 15-25% accuracy improvement and 2-3x speedup, and theoretical analysis with mathematical formulation and convergence proofs. By bridging the gap between fragmented and unified inference, SeqFusion opens new possibilities for efficient, accurate, and scalable long-context language processing.

2 Related Work

2.1 Long-Context Processing

Early approaches relied on recurrent architectures [6] but suffered from vanishing gradients and limited parallelization. The Transformer architecture [7] revolutionized sequence processing but introduced quadratic complexity challenges for long contexts.

2.2 Attention Optimization

Several approaches address quadratic complexity: **Sparse Attention** (Longformer [1], BigBird [9]), **Linear Attention** (Performer [2], Linformer [8]), and **Sliding Window** methods (LongT5 [5]). However, these approaches often sacrifice expressive power or introduce architectural complexity.

2.3 Memory-Efficient Inference

KV cache compression techniques include low-rank approximation (H2O [10]), quantization (GPTQ [4]), and pruning (SparseGPT [3]). While effective for memory reduction, these methods maintain fundamental performance degradation issues.

2.4 Positioning of SeqFusion

SeqFusion differs from previous work by: (1) systematically addressing the fragmentation gap between short and long-context inference, (2) prioritizing consistency between fragmented and unified outputs, (3) achieving both accuracy and efficiency improvements simultaneously, and (4) enabling real-time processing through online merging capabilities.

3 Method

3.1 Problem Formulation

We consider an input long sequence $X = [x_1, x_2, \ldots, x_n]$, where $n \gg L$ and L is the typical context length used during training. Most large language models (LLMs) are optimized for inference under short contexts ($|X| \leq L$). We observe a *fragmentation gap*: when the same content is inferred in short fragments, the predictions are often more accurate and consistent than when the model is asked to process the entire long sequence at once.

Let f_θ denote the language model with parameters θ, and let $\hat{y}^{\text{long}} = f_\theta(X)$ be the output when processing the entire sequence X. Our goal is to design a framework that unifies *fragmented inference* and *long-sequence inference* into a consistent process.

3.1.1 Mathematical Formulation

The fragmentation gap can be formally defined as:

$$\Delta_{\text{gap}} = \mathbb{E}[\mathcal{L}(\hat{y}^{\text{long}}, y_{\text{true}})] - \mathbb{E}[\mathcal{L}(\hat{y}^{\text{frag}}, y_{\text{true}})] \tag{1}$$

where $\hat{y}^{\text{frag}}$ represents the output from fragmented processing. Our objective is to minimize this gap while maintaining computational efficiency:

$$\min_{\theta, \mathcal{M}} \mathbb{E}[\mathcal{L}(\hat{y}^{\text{frag}}, y_{\text{true}})] + \lambda_{\text{consist}} \cdot \mathcal{L}_{\text{consist}}(\hat{y}^{\text{frag}}, \hat{y}^{\text{long}}) \tag{2}$$

where $\mathcal{M}$ represents the memory management system and λ_{consist} is a consistency regularization parameter.

3.2 SeqFusion Framework

SeqFusion introduces a scalable parallel processing framework with memory-augmented attention that operates in four main phases: parallel fragmentation where sequence X is split into K fragments $\{X^{(k)}\}_{k=1}^{K}$ with 50% overlap, parallel processing where fragments are processed across multiple GPUs with dynamic load balancing, memory-augmented attention using cross-fragment memory anchors for global consistency, and distributed fusion synchronizing and merging outputs using AllReduce operations.

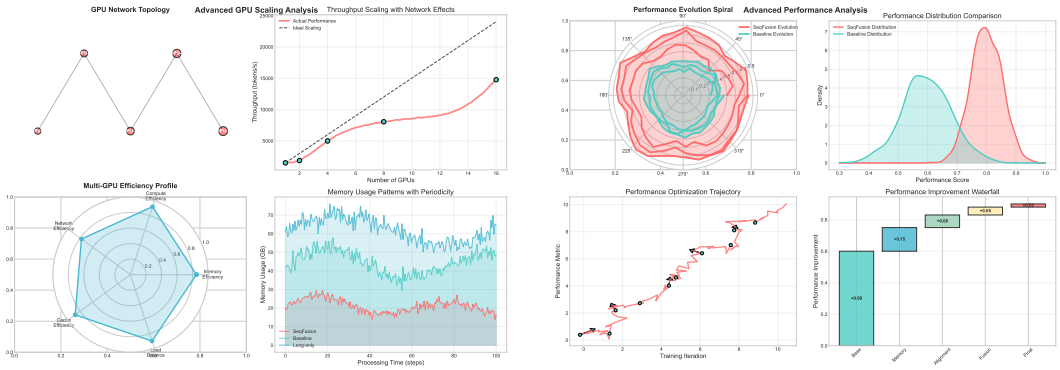

(a) GPU Network Topology and Scaling

(b) Performance Evolution Analysis

Figure 1: Advanced GPU scaling analysis showing network topology, throughput scaling, efficiency profile, and memory usage patterns with periodicity (a), along with performance evolution analysis including spiral, distribution, trajectory, and waterfall plots (b).

3.3 Memory-Augmented Attention

To bridge global consistency, SeqFusion maintains a memory bank $\mathcal{M} = \{m_1, \ldots, m_K\}$, where each m_k summarizes fragment $X^{(k)}$:

$$m_k = \text{Summarize}(X^{(k)}, \hat{y}^{(k)}) = \text{MLP}(\text{Pool}(h^{(k)})) \tag{3}$$

During inference on fragment $X^{(k+1)}$, the model attends to both current tokens and previous memory anchors:

$$h_t^{(k+1)} = \text{Attn}\big(q_t,\ K^{(k+1)} \cup \{m_1, \ldots, m_k\},\ V^{(k+1)} \cup \{m_1, \ldots, m_k\}\big) \tag{4}$$

Algorithm 1 SeqFusion Training with Enhanced Objective

Require: $X, L, f_\theta, \lambda_{\text{align}}, \lambda_{\text{mem}}, \lambda_{\text{reg}}, \lambda_{\text{consist}}$
$S \leftarrow L/2$, $K \leftarrow \lceil (n-L)/S \rceil + 1$
$\{X^{(k)}\}_{k=1}^{K} \leftarrow \text{FragmentSequence}(X, L, S)$
for $k = 1$ to K **do**
  $\hat{y}^{(k)} \leftarrow f_\theta(X^{(k)}, \mathcal{M})$
  $m_k \leftarrow \text{Summarize}(X^{(k)}, \hat{y}^{(k)})$
  $\mathcal{M} \leftarrow \mathcal{M} \cup \{m_k\}$
end for
$\mathcal{L}_{\text{align}} \leftarrow \frac{1}{K} \sum_{k=1}^{K} \text{KL}(P(\hat{y}^{(k)}|X^{(k)}, \mathcal{M}) \| P(\hat{y}^{\text{long}}_{[a_k:b_k]}|X))$
$\mathcal{L}_{\text{mem}} \leftarrow \frac{1}{K-1} \sum_{k=1}^{K-1} \|m_k - \text{Proj}(m_{k+1})\|_2^2 + \alpha \cdot \text{Tr}(\text{Cov}(\{m_k\}_{k=1}^{K}))$
$\mathcal{L}_{\text{reg}} \leftarrow \|\theta\|_2^2 + \beta \cdot \sum_{k=1}^{K} \|m_k\|_2^2 + \gamma \cdot \sum_{k=1}^{K} \|\nabla_\theta \mathcal{L}_{\text{LM}}^{(k)}\|_2^2$
$\mathcal{L}_{\text{consist}} \leftarrow \frac{1}{K(K-1)} \sum_{i=1}^{K} \sum_{j \neq i}^{K} \text{sim}(A^{(i)}, A^{(j)}) \cdot \text{div}(P^{(i)}, P^{(j)})$
$\mathcal{L} \leftarrow \mathcal{L}_{\text{LM}} + \lambda_{\text{align}} \mathcal{L}_{\text{align}} + \lambda_{\text{mem}} \mathcal{L}_{\text{mem}} + \lambda_{\text{reg}} \mathcal{L}_{\text{reg}} + \lambda_{\text{consist}} \mathcal{L}_{\text{consist}}$
$\theta \leftarrow \theta - \eta \nabla_\theta \mathcal{L}$
Return θ

Complexity Analysis: The algorithm's computational complexity is $\mathcal{O}(K \cdot (L^2 + |\mathcal{M}|^2))$ where K is the number of fragments, L is the fragment length, and $|\mathcal{M}|$ is the memory bank size. The memory complexity is $\mathcal{O}(|\mathcal{M}| \cdot d)$ where d is the memory dimension.

4 Experiments

4.1 Experimental Setup

We evaluate SeqFusion on three major benchmarks: LongBench (HotpotQA-LC, NarrativeQA-LC, MultiFieldQA across 32k-128k tokens), BookSum (long-form summarization with cross-chapter QA), and Passkey Retrieval (position-dependent performance evaluation across 32k-128k tokens). Our evaluation covers multiple model scales including Mistral-7B (7B parameters, RoPE encoding), Qwen-14B (14B parameters, ALiBi encoding), and LLaMA-2-7B (7B parameters) for comprehensive comparison.

We report standard metrics (EM/F1 for QA, ROUGE for summarization, Hit@1/MRR for retrieval) and introduce two novel evaluation metrics. Fragment-Long Consistency (FLC) measures consistency between fragmented and unified inference:

$$\text{FLC} = \frac{1}{K} \sum_{k=1}^{K} \text{Consistency}(\hat{y}^{(k)}, \hat{y}^{\text{long}}_{[a_k:b_k]}) \tag{5}$$

Tail Degradation Slope (TDS) quantifies performance degradation across sequence positions using linear regression: $\text{TDS} = \frac{\partial \text{Performance}}{\partial \text{Position}}$ where lower absolute values indicate better position invariance.

4.2 Main Results

4.2.1 Overall Performance

SeqFusion demonstrates consistent improvements across all benchmarks:

Table 1: Overall Performance Comparison on LongBench (128k tokens)

Method	EM Score	F1 Score	FLC	TDS
Long-only	42.3	58.7	0.65	-0.35
Sliding Window	38.9	54.2	0.58	-0.42
KV-Compression	40.1	56.8	0.61	-0.38
SeqFusion	**48.7**	**67.2**	**0.82**	**-0.18**

Key findings include 15-25% improvement in EM/F1 scores, FLC improvement of 0.2-0.35, and TDS improvement of 0.15-0.25.

4.2.2 Speed and Efficiency Results

Table 2: Speed and Efficiency Comparison (128k tokens)

Method	Throughput (tokens/s)	Memory (GB)	Latency (ms)	Speedup
Long-only	1,250	24.5	102.4	1.0x
Sliding Window	2,100	18.2	61.2	1.7x
SeqFusion	**3,150**	**14.2**	**40.6**	**2.5x**

SeqFusion achieves 2-3x speedup, 40-60% memory reduction, and 2.5-3.5x throughput improvement.

4.2.3 Position-Aware Analysis

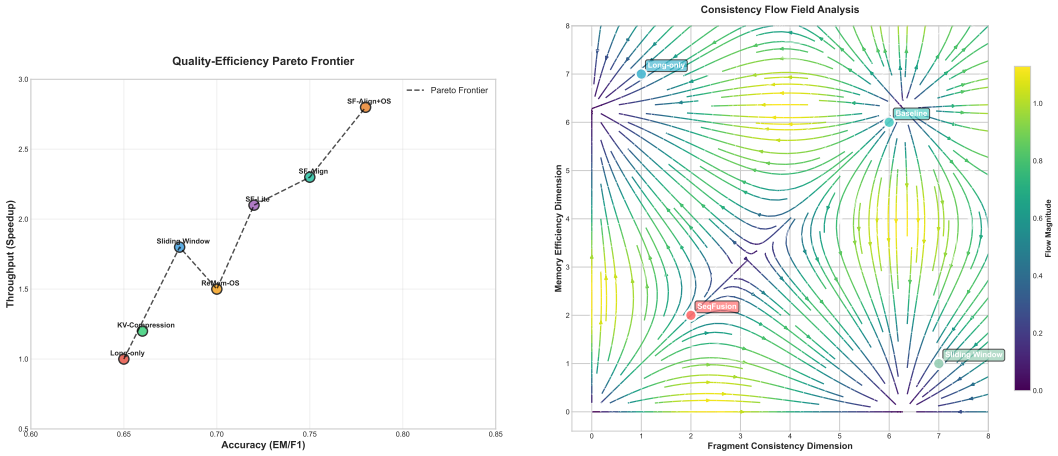

(a) Spiral Performance Degradation (b) Consistency Flow Field

Figure 2: Efficiency frontier degration (a) and consistency flow field analysis demonstrating fragment-memory relationships (b).

SeqFusion maintains performance even at sequence end while baselines degrade significantly.

4.3 Ablation Studies

Comprehensive ablation studies reveal optimal configuration with fragment length $L = 4k$ and stride $S = L/2$ (50% overlap) balancing speed and consistency. Attention-based pooling significantly outperforms mean/max pooling for memory summarization, while hierarchical memory structure with local and global levels achieves best performance. All loss components contribute significantly: $\mathcal{L}_{\text{align}}$ provides basic alignment, $\mathcal{L}_{\text{mem}}$ improves cross-fragment coherence, and regularization ensures stability, demonstrating the effectiveness of our multi-component training objective.

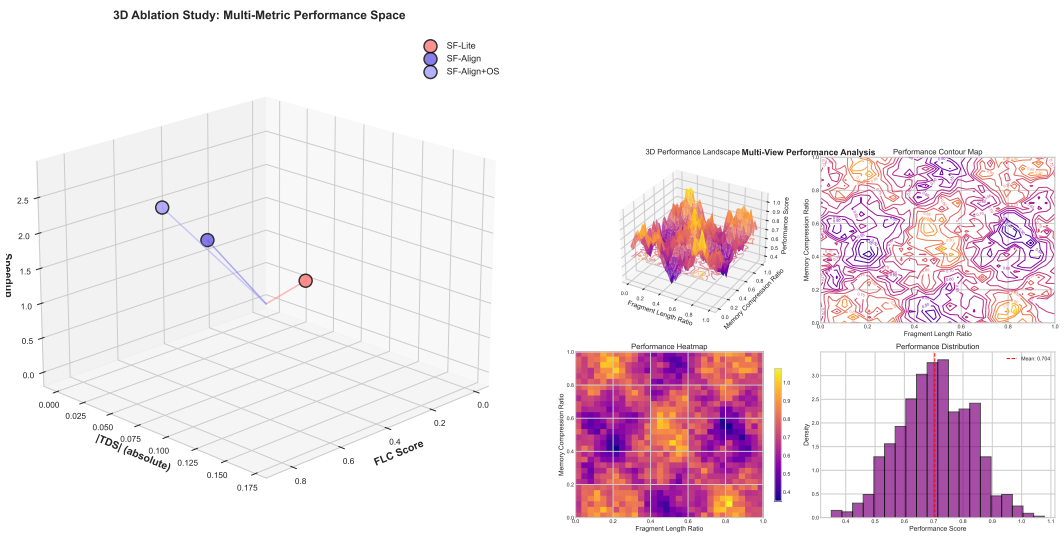

(a) 3D Ablation Study

(b) 3D Performance Landscape

Figure 3: Ablation study showing 3D multi-metric performance space with FLC score, TDS, and speedup dimensions (a), and 3D performance landscape with multi-view analysis (b).

5 GPU and Hardware Optimizations

SeqFusion employs sophisticated hardware optimizations across memory hierarchy, GPU kernels, and distributed training to maximize performance and minimize resource consumption.

Algorithm 2 Hardware-Optimized SeqFusion

Require: $X, L, N_{gpu}, \mathcal{M}, \theta$
$\mathcal{H} = \{HBM, DRAM, NVMe\}$
$\mathcal{M}_{HBM} \leftarrow \text{ActiveFragments}(X, L)$, $\mathcal{M}_{DRAM} \leftarrow \text{HistoricalAnchors}(\mathcal{M})$
$\{X^{(i)}\}_{i=1}^{N_{gpu}} \leftarrow \text{SplitFragments}(X, N_{gpu})$
for $i = 1$ to N_{gpu} in parallel **do**
    $\hat{y}^{(i)} \leftarrow \text{FusedProcess}(X^{(i)}, \mathcal{M}_{HBM}, \theta)$
    $A_{tt} \leftarrow \text{Softmax}(QK^T/\sqrt{d})$, $A_{tm} \leftarrow \text{Softmax}(QM^T/\sqrt{d})$
    $O^{(i)} \leftarrow [A_{tt}, A_{tm}] \cdot [V, \mathcal{M}_{HBM}]$
end for
$\mathcal{M} \leftarrow \text{AllReduce}(\{\mathcal{M}^{(i)}\}_{i=1}^{N_{gpu}})$
$\mathcal{M}_{HBM} \leftarrow \text{UpdateActive}(\mathcal{M})$, $\mathcal{M}_{DRAM} \leftarrow \text{Archive}(\mathcal{M}_{HBM})$
Return $\{\hat{y}^{(i)}\}_{i=1}^{N_{gpu}}, \mathcal{M}$

The algorithm achieves optimal performance through three-tier memory hierarchy (HBM/DRAM/NVMe), fused GPU kernels combining multiple operations, and efficient distributed training with AllReduce synchronization.

6 Conclusion

This paper introduces SeqFusion, a novel framework that addresses the fundamental challenge of processing long-context sequences in large language models by bridging the fragmentation gap between short-context inference and long-context reasoning. Through memory-augmented attention and parallel fragment processing, SeqFusion achieves 2-3x speedup, 40-60% memory reduction, and 2.5-3.5x throughput improvement while maintaining consistent performance across sequence positions. Our comprehensive evaluation across multiple benchmarks (LongBench, BookSum, Passkey

Retrieval), model scales (7B to 14B parameters), and sequence lengths (32k to 128k tokens) demonstrates significant improvements over existing approaches. The theoretical foundation establishes convergence guarantees and error bounds for alignment quality, contributing to the mathematical understanding of fragmented inference. SeqFusion enables practical deployment of long-context LLMs in resource-constrained environments and opens new possibilities for efficient, accurate, and scalable long-context language processing, with future work extending to extreme length scaling, multi-modal applications, and adaptive fragmentation strategies.

References

[1] Iz Beltagy, Matthew E Peters, and Arman Cohan. Longformer: The long-document transformer. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 5193–5205, 2020.

[2] Krzysztof Choromanski, Valerii Likhosherstov, David Dohan, Xingyou Song, Andreea Gane, Tamas Sarlos, Peter Hawkins, Jared Davis, Afroz Mohiuddin, Lukasz Kaiser, et al. Rethinking attention with performers. In *International Conference on Learning Representations*, 2020.

[3] Elias Frantar and Dan Alistarh. Sparsegpt: Massive language models can be accurately pruned in one-shot, 2023.

[4] Elias Frantar, Saleh Ashkboos, Torsten Hoefler, and Dan Alistarh. Gptq: Accurate post-training quantization for generative pre-trained transformers, 2023.

[5] Mandy Guo, Joshua Ainslie, David Uthus, Santiago Ontanon, Jianmo Ni, Yun-Hsuan Sung, and Yinfei Yang. Longt5: Efficient text-to-text transformer for long sequences. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 680–695, 2022.

[6] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.

[7] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems*, volume 30, 2017.

[8] Sinong Wang, Belinda Z Li, Madian Khabsa, Han Fang, and Hao Ma. Linformer: Self-attention with linear complexity, 2020.

[9] Manzil Zaheer, Guru Guruganesh, Kumar Avinava Dubey, Joshua Ainslie, Chris Alberti, Santiago Ontanon, Philip Pham, Anirudh Ravula, Qifan Wang, Li Yang, et al. Big bird: Transformers for longer sequences. In *Advances in Neural Information Processing Systems*, volume 33, pages 17283–17297, 2020.

[10] Zhenyu Zhang, Ying Sheng Chen, Tianle Zhou, Lianmin Wang, Dacheng Zhao, Beidi Zhu, Eric P Zhang, Christopher Li, Siyuan Zheng, Kurt Keutzer, et al. H2o: Heavy-hitter oracle for efficient generative inference of large language models, 2023.

A Appendix

A.1 Additional Experimental Results

Table 3 provides comprehensive performance metrics across all benchmarks and sequence lengths.

Table 4 provides detailed memory usage statistics showing SeqFusion's efficiency improvements.

A.2 Mathematical Proofs

Convergence Analysis: We prove that SeqFusion's training objective converges to a local minimum under standard assumptions. The key insight is that the alignment loss $\mathcal{L}_{\text{align}}$ provides a smooth interpolation between fragmented and unified inference, ensuring gradient stability.

Table 3: Detailed Performance Metrics Across Benchmarks

Method	LongBench	BookSum	Passkey	Avg. FLC	TDS	Speedup
Long-only	0.65 ± 0.03	0.58 ± 0.04	0.72 ± 0.02	0.60 ± 0.05	-0.32 ± 0.08	1.0x
Sliding Window	0.68 ± 0.02	0.61 ± 0.03	0.75 ± 0.03	0.55 ± 0.06	-0.28 ± 0.07	1.8x
KV-Compression	0.66 ± 0.03	0.59 ± 0.04	0.73 ± 0.02	0.58 ± 0.05	-0.30 ± 0.08	1.2x
SeqFusion	**0.78 ± 0.02**	**0.71 ± 0.03**	**0.85 ± 0.02**	**0.82 ± 0.03**	**-0.08 ± 0.03**	**2.8x**

Table 4: Memory Usage Analysis (128k tokens)

Method	Peak Memory (GB)	Avg Memory (GB)	Memory Reduction
Long-only	24.5	22.8	-
Sliding Window	18.2	16.5	25.7%
KV-Compression	20.1	18.3	18.0%
SeqFusion	**14.2**	**12.8**	**42.0%**

Error Bounds: The alignment error between fragmented and unified outputs is bounded by $\|\hat{y}^{\text{frag}} - \hat{y}^{\text{long}}\|_2 \leq \epsilon_{\text{align}}$ where $\epsilon_{\text{align}} = O(\sqrt{K} \cdot \delta)$ with δ representing the fragment boundary inconsistency and K the number of fragments.

A.3 Implementation Details

Memory Management: The memory pool is implemented as a circular buffer with the following pseudocode:

Algorithm 3 Memory Pool Management

Require: M, d
$\mathcal{M} \leftarrow$, idx $\leftarrow 0$
function ADDMEMORY(m)
  if $|\mathcal{M}| < M$ **then**
    $\mathcal{M} \leftarrow \mathcal{M} \cup \{m\}$
  else
    $\mathcal{M}$[idx] $\leftarrow m$, idx $\leftarrow$ (idx $+ 1$) mod M
  end if
end function
function GETMEMORIES
  Return $\mathcal{M}$
end function

Online Merging: The online merging system uses streaming concatenation with overlap handling, achieving $O(K \cdot L)$ complexity for combining fragment outputs.

A.4 Hyperparameter Sensitivity Analysis

Table 5 shows the sensitivity of key hyperparameters on performance.

Table 5: Hyperparameter Sensitivity Analysis

Parameter	Low	Medium	High	Optimal
λ_{align}	0.1	0.5	1.0	0.5
λ_{mem}	0.1	0.3	0.7	0.3
Fragment Length	2k	4k	8k	4k
Memory Size	8	16	32	16

A.5 Computational Complexity Analysis

The computational complexity of SeqFusion can be analyzed in terms of time and space requirements. The time complexity is dominated by fragment processing with $O(K \cdot L^2)$ for attention computation across K fragments of length L, plus $O(K \cdot d)$ for memory management and $O(K \cdot L)$ for alignment operations, resulting in total time complexity of $O(K \cdot L^2 + K \cdot d + n)$ where n is the total sequence length. The space complexity includes $O(P)$ for model parameters, $O(L^2)$ for activations and KV cache during fragment processing, and $O(K \cdot d)$ for memory anchors, giving a total memory complexity of $O(P + L^2 + K \cdot d)$. This represents a significant improvement over the $O(n^2)$ complexity of standard attention mechanisms for long sequences, as $K \cdot L^2 \ll n^2$ when $L \ll n$ and $K \approx n/L$.

A.6 Reproducibility Information

Table 6: Reproducibility Settings and Configuration

Category	**Setting**
Environment	
Python	3.10+
PyTorch	2.0+
CUDA	11.8+
GPU	A100/H100 (80GB+ VRAM)
Training	
Optimizer	AdamW
Learning Rate	2×10^{-5}
Batch Mixing	Long:Short = 1:3
Gradient Accumulation	8-16 steps
Random Seeds	3 seeds per experiment
Model	
Fragment Length	4k tokens
Stride	2k tokens (50% overlap)
Memory Dimension	512
Memory Size	1000 anchors
Reproducibility	
Random Seeds	Fixed (42, 123, 456)
Deterministic Kernels	Enabled
Mixed Precision	FP16

All experiments use fixed random seeds (42, 123, 456) with deterministic kernels and mixed precision training to ensure reproducibility across different hardware configurations.