End-to-End Context Compression at Scale

Ang Li [♠▲]* **Sean McLeish** [♥]* **Haozhe Chen** [♣]* **Nimit Kalra** [♦]
Zaiqian Chen [♦] **Artem Gazizov** [★] **Venkata Anoop Suhas Kumar Morisetty** [♦]
Bhavya Kailkhura [●] **Harshitha Menon** [●] **Zhuang Liu** [♣] **Brian R. Bartoldson** [●]
Tom Goldstein [♥] **Sanae Lotfi** [■1] **Micah Goldblum** [♦‡] **Pavel Izmailov** [♠‡]

[♠]New York University [▲]Modal Labs [♥]University of Maryland [♣]Princeton University
[♦]Columbia University [★]Harvard University [●]Lawrence Livermore National Laboratory
[■]FAIR at Meta

*Equal contribution, correspondence to al6843@nyu.edu, smcleish@umd.edu, hc5019@princeton.edu

‡Equal advising, correspondence to micah.g@columbia.edu, pi390@nyu.edu

arXiv:2606.09659v2 [cs.CL] 18 Oct 2026

Long-context language model inference is bottlenecked by memory, as the KV cache grows with context length. Recent techniques to compress the KV cache fall short: they either degrade model quality substantially or require considerable time and compute to compress a single long prompt. Furthermore, many methods require the input to fit within the target model's context window, and are generally incompatible with modern production inference engines. Encoder-decoder compressors, which map a long token sequence to a shorter sequence of latent embeddings consumed by a decoder, are an appealing alternative in principle. However, existing approaches are not competitive with KV cache compression on the accuracy-efficiency frontier. In this work, we revisit encoder-decoder compression and close this gap. We first perform an architecture search, pre-training many variants from scratch to determine how best to design and train encoder-decoder compressors. Guided by our findings, we continually pre-train a family of 0.6B-encoder, 4B-decoder models on nearly 400B tokens each, at compression ratios of 1:4, 1:8, and 1:16. We introduce *Latent Context Language Models* (LCLMs), a family of compressors that improve the Pareto frontier across general-task performance, compression speed, and peak memory usage. LCLMs also unlock new efficient downstream applications: they serve as backbones for long-horizon agents that skim a compressed long context and adaptively expand relevant segments on demand, and they can compact the model's own generated reasoning into latent memory during inference, which is nearly lossless on the full MATH test set while being highly efficient, reducing peak KV usage by up to 7.9×.

🤗 **Models**: huggingface.co/latent-context
Code: github.com/LeonLixyz/LCLM

1 Introduction

Reasoning over long contexts is a crucial capability for state-of-the-art Large Language Models (LLMs), as it enables them to parse long documents, engage in multi-turn conversations, and perform long-horizon agentic tasks. However, in production systems, the input context, working horizon, and memory can grow to millions of tokens, making inference increasingly constrained by memory and latency due to the growth of the KV cache (Hooper et al., 2024). Even when inputs fit within the model's maximum context window, models often struggle to reliably use information distributed across long contexts (Liu et al., 2024b; An et al., 2025). As a result, context and memory management have emerged as critical systems and modeling challenges for deploying long-horizon LLM systems.

A natural response is to reduce the size of the KV cache directly. KV cache compression methods reduce the memory footprint of cached activations by evicting cache entries (Li et al., 2024c; Devoto et al., 2025; Kim et al., 2025; Zweiger et al., 2026) or training a compact cache offline (Eyuboglu et al., 2025). However, they have important limitations for general long-context inference. Many methods still require the full context to be prefilled before compression, and some require substantial additional time or memory beyond the original prefill. Query-dependent methods produce caches that are difficult to reuse across turns (Li et al., 2024b), while methods that evict caches non-uniformly across attention heads and layers (Kim et al., 2025) are difficult to integrate with modern inference engines that assume a shared sequence length across the cache.

[1]Served solely in an advisory role for this paper. All model training and data processing were conducted entirely outside of Meta.

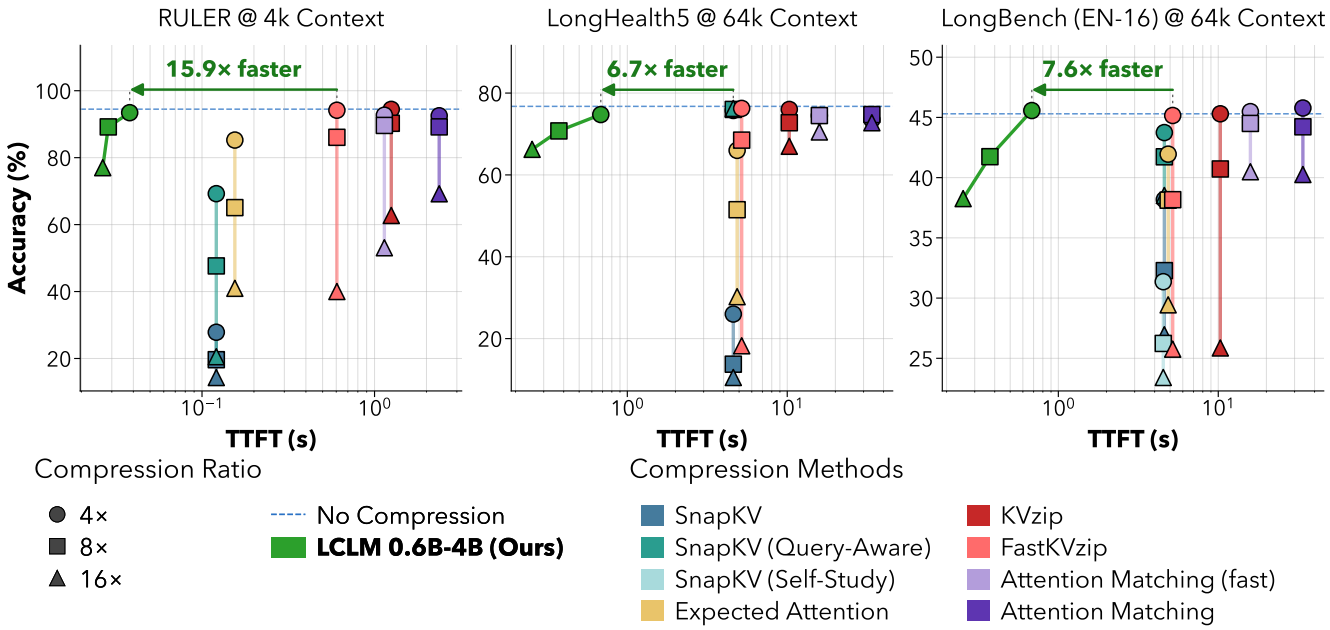


Figure 1 Our Latent Context Language Models achieve high quality compression while being fast and memory efficient. We show accuracy vs. TTFT (time to first token) per sample on a single H200 for RULER (Hsieh et al., 2024) at 4k tokens, and for LongHealth5 (Adams et al., 2025) and LongBench (Bai et al., 2024) at 64k tokens. All methods are based on the same decoder `Qwen3-4B-Instruct-2507` (Yang et al., 2025). We see that our models lie on a new Pareto frontier, compressing contexts much faster while maintaining high accuracy, especially for high compression ratios. KV cache compression baselines appear as vertical lines because the context is prefilled the same way regardless of compression ratio, and the eviction operation is much faster than the prefill time. Arrows mark the speedup of the LCLM at 4× compression over the fastest KV cache method at 4× compression with accuracy within 1 point of it, so the matched baseline, and hence the speedup, can differ between benchmarks of the same length.

Encoder-decoder soft-token methods (e.g. Lin et al., 2025; Ge et al., 2023; Chevalier et al., 2023; Liao et al., 2025; Yen et al., 2024) offer a promising way to address the limitations of KV cache compression. Rather than manipulating the KV cache directly, these models encode the raw input tokens into a much shorter sequence of continuous embeddings, which are then provided to the LLM in place of the original context. In principle, this approach is highly attractive: compression is parallelizable, supported by standard LLM inference engines, and works with standard LLM decoding while extending the decoder far beyond its native context length.

However, existing soft-token compression methods typically face a trade-off: they either substantially degrade the capabilities of the base language model or depend on domain-specific training to remain effective. This motivates a fundamental question: **can we train a general, task-agnostic compressor that preserves a model's original capabilities?** To this end, we introduce *Latent Context Language Models* (LCLMs), a family of encoder-decoder compressors trained end-to-end at scale. We initialize both the encoder and decoder with pre-trained models and jointly train them on a total of nearly 400B tokens for the combined encoder-decoder. Crucially, we demonstrate that LCLMs can preserve decoder performance while reducing memory used for input context.

Our contributions are summarized as follows:

1. **Training recipe:** We develop a fine-grained training recipe for encoder-decoder context compression. Accordingly, we curate continual pre-training and finetuning data for training compressors and optimize training details across multiple training stages. See Section 4.
2. **Comprehensive architectural search:** We conduct a large-scale architecture search to identify optimal design choices for encoder-decoder compression models in a controlled setting. See Section 5 and Figure 3.
3. **New memory/speed-performance frontier:** We train a family of models across different compression ratios at scale. Our experiments identify a Pareto frontier that effectively balances memory, latency, and downstream accuracy. We open-source this suite of models and all code. See Section 6 and Figures 1, 4 and 5.
4. **New capabilities from compression:** We create an agentic system with a high compression ratio, where

the agent can select which compressed chunk to expand, substantially enhancing the model's performance on challenging needle-in-a-haystack tasks. We also use the encoder to compress the model's own generated tokens during decoding, so long chains of thought are kept in a compact latent memory; on MATH, this is nearly lossless while reducing KV usage several-fold. See Section 7 and Figures 6 and 7.

2 Related Work

Context compression approaches generally fall into three categories: hard-token, soft-token, and KV cache compression. We defer the discussion of hard-token approaches to Appendix A, since they generally underperform compared to the latter two.

KV cache compression. KV cache compression methods evict entries from the KV cache. Most methods rely on hand-crafted or learned heuristics to select which entries to drop, while others train the model to merge or sparsify its own cache (Nawrot et al., 2024; Łańcucki et al., 2026).

Prompt-agnostic methods (Kim et al., 2025; Devoto et al., 2025; Zweiger et al., 2026) prune the context without knowledge of the query, whereas prompt-dependent methods such as SnapKV (Li et al., 2024c) require explicit context-prompt pairs and yield query-specific caches that limit multi-turn applicability (Li et al., 2024b; Kim et al., 2025). An alternative is to anticipate future queries at compression time: Cartridges (Eyuboglu et al., 2025) distills a fixed-size KV cache per corpus. However, this comes at a substantial training cost per corpus, with an ICL-quality cache requiring approximately 30 minutes on an 8×H100 node for an 8B model.

Although KV cache compression methods are well-studied, they are not widely adopted in inference engines such as vLLM (Kwon et al., 2023) or SGLang (Zheng et al., 2024). Compression methods such as KVzip allocate eviction budgets non-uniformly across attention heads and layers; therefore, they cannot reduce the sequence-length dimension of the KV cache. Instead, KV cache compression methods mask evicted positions in the attention computation, forfeiting the memory and throughput benefits of these paged-attention engines. Some methods further require prefill passes that are two to three times longer than the original input (Kim et al., 2025). We note that soft token compression methods are fully compatible with these open-source inference frameworks.

Soft-token compression. Despite promising results, most prior soft-token methods rely on offline preprocessing and do not convincingly preserve the base model's broad in-context behavior. Prior work is mostly evaluated on domain- or task-specific finetuning that is not transferable across tasks (Tang et al., 2025; Feldman and Artzi, 2025; Dai et al., 2025; Lin et al., 2025; Cheng et al., 2024). To the best of our knowledge, these methods are not evaluated on long-context benchmarks with information-dense tasks that require fine-grained details throughout the whole context, such as RULER (Hsieh et al., 2024), to demonstrate complex long-context understanding.

In Section 5, we discuss the architectural decisions of prior work on soft-token compression as part of our architectural search and analysis. Our work shows that, with large-scale staged training, a single online soft-token compressor can handle heterogeneous inputs while closely matching uncompressed behavior.

3 The Latent Context Language Model Architecture

An LCLM consists of an encoder with a pooling method that maps token chunks to soft tokens, an adapter to project to the decoder's embedding dimension, and a decoder that consumes the soft tokens as latent context. Given a sequence of input tokens $x_{1:T}$ and a compression ratio N, the encoder maps each contiguous block of N input tokens to a single latent token. Let W denote the encoder window size, i.e., the number of input tokens processed by the encoder in one forward pass. We split the input sequence into $I = \lceil T/W \rceil$ encoder windows, each containing at most W tokens:

$$w_i = x_{(i-1)W+1\,:\,\min(iW,T)}, \qquad i = 1, \ldots, I. \tag{1}$$

For each window w_i, the encoder computes the final hidden states $h^{(i)}_{1:|w_i|}$ for each input token in the window. A pooling operator then aggregates these hidden states into $M_i = \left\lceil \frac{|w_i|}{N} \right\rceil$ latent tokens $z^{(i)}_{1:M_i}$. When $W = N$, each encoder window contains a single compression block, yielding one latent token per forward pass. When $W = T$, the encoder processes the entire input sequence in one forward pass and produces $\lceil T/N \rceil$ latent tokens. We concatenate the latent tokens from all encoder windows to obtain the full compressed latent sequence $z_{1:M}$, where M is the total number of latent tokens after compression.

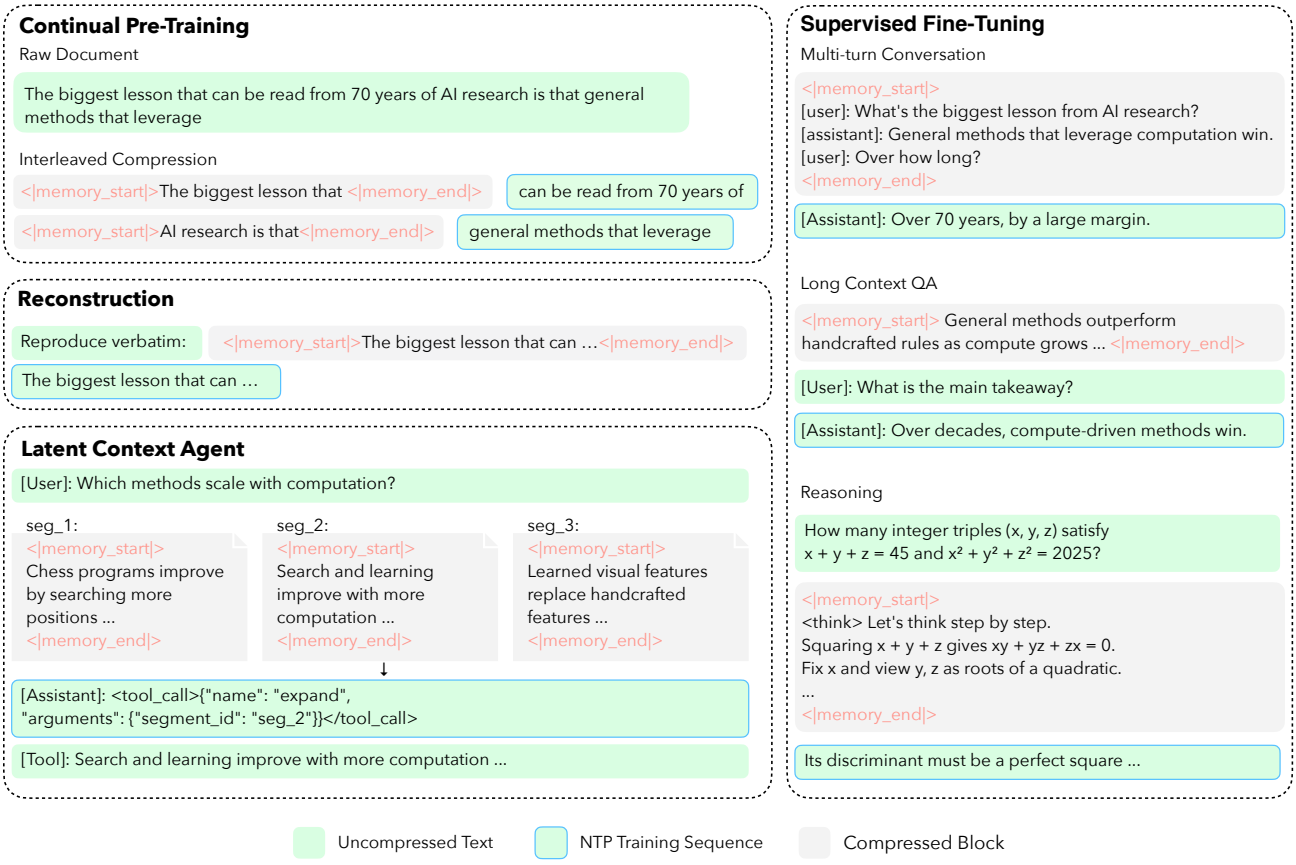


Figure 2 Examples of the data used to train LCLMs. Continual pre-training interleaves compressed and uncompressed spans, and reconstruction recovers source text from latents. SFT includes conversations, long-context QA, reasoning, and agentic trajectories. Reasoning questions remain raw while selected spans of the reference chain of thought are compressed; the remaining continuation is supervised. Native tool trajectories remain raw, while expansion examples compress source observations and retain raw tool calls and returned text. Blue outlines mark supervised decoder tokens.

Because the encoder and decoder may have different hidden dimensions, an adapter $a(\cdot)$ projects the compressed latent sequence $z_{1:M}$ from the encoder hidden dimension into the decoder hidden dimension, producing latent tokens $s_{1:M}$. The decoder then consumes these projected latent tokens in place of the corresponding input tokens.

4 Training Recipe

We now describe our multi-stage recipe and data for training. In contrast to prior context-compression work that trains specialized models on small-scale in-domain datasets (Tang et al., 2025; Feldman and Artzi, 2025; Pilchen et al., 2025; Cheng et al., 2024; Liao et al., 2025), our goal is to preserve the strong performance of a powerful LLM across downstream tasks. To this end, we curate three types of data: continual pre-training data, Supervised Fine-Tuning (SFT) data, and auxiliary reconstruction data.

4.1 Training Data

Continual pre-training dataset. We construct a carefully curated high-quality pre-training mixture covering Common Crawl text, code, science and reasoning, long-context data, and instruction-style data. We partition each sequence into multiple segments and alternate between compressed segments and standard token segments. We then compute the next-token prediction loss on the uncompressed tokens *only*. This interleaved format differs from the simple first-half compressed, second-half trainable split used in prior work (Lin et al., 2025; Ge et al., 2023; Chevalier et al., 2023; Yen et al., 2024). By distributing compressed spans throughout the sequence, the model learns

to condition on latent context at multiple positions rather than only at the beginning of the context. We provide a visual example of this data format in Figure 2. The exact pre-training mixture is described in Appendix B.1, with additional details on compression formatting in Appendix B.4.

SFT data. To recover performance after continual pre-training (Gao et al., 2025) and teach the model to use latent context, we post-train with supervised fine-tuning on compressed inputs. Our SFT mixture spans four components: (i) reasoning, (ii) long-context instruction following, (iii) general instruction following and multi-turn conversation, and (iv) agentic trajectories. For long-context and instruction data, we compress the documents and prompts. For reasoning, we instead keep the question uncompressed and compress spans of the reference chain of thought, so the model learns to continue reasoning from compressed spans. To preserve tool use under compressed inputs, we also include public tool-calling, search, and terminal trajectories (Blakeman et al., 2025; Raoof et al., 2026), kept uncompressed, and 0.15M synthetic trajectories (under 1% of the mixture) in which the model calls an `expand` tool to read back compressed legal and contract documents; these teach the tool interface of Section 7.1 but share no data with its RULER tasks. Because some original completions were generated by outdated models, we relabel a subset of the SFT data using `Qwen3-30B-A3B-Instruct-2507` and `Qwen3-235B-A22B-Instruct-2507`, released under the Apache-2.0 license. The exact mixture of 21.5M examples is reported in Appendix B.3.

Auxiliary reconstruction data. To accelerate training and encourage the latent representations to preserve fine-grained information, we introduce an auxiliary reconstruction task. Given a sampled document, we compress the text and ask the model to repeat the original content. To increase diversity, we include a variety of sources spanning code, text, long documents, math, and LaTeX, as described in Appendix B.2. To reduce prompt overfitting, we use a bank of 100 prompt templates per source. This auxiliary task encourages the compressor not only to support semantic understanding but also to retain fine-grained details needed for tasks such as exact retrieval. We include auxiliary reconstruction data in both continual pre-training and SFT.

4.2 Training Stages

We adopt a multi-stage training recipe to improve stability, preserve the original model's capabilities, and mitigate catastrophic forgetting. The stages progressively increase the number of trainable components, starting with adapter and encoder warm up, followed by end-to-end continual pre-training and SFT.

- **Stage 0: Adapter warmup.** We freeze both the encoder and decoder and update only the adapter.
- **Stage 1: Encoder training.** We unfreeze the encoder while keeping the decoder frozen.
- **Stage 2: End-to-end continual pre-training.** We unfreeze the decoder with a small learning rate.
- **Stage 3: Supervised fine-tuning.** We train on the SFT mixture and increase the decoder learning rate.

This staged procedure is analogous to Vision-Language Model (VLM) training pipelines that first align pre-trained representations across modalities (Liu et al., 2023a; Tong et al., 2024) and then fine-tune for downstream tasks: stages 0 and 1 play the role of projector alignment, stage 2 corresponds to interleaved continual pre-training, and stage 3 to instruction tuning. For stages 0, 1, and 2, we uniformly sample from the continual pre-training mixture; for stage 3, we train on the SFT mixture. In early experiments, we consider directly training the full model end-to-end from the beginning, but find that this approach underperforms the multi-stage recipe. At the beginning of training, the decoder is not accustomed to taking in the output of the embedding model, so both encoder and decoder parameter gradients can be large, leading to model degradation. The staged approach, where we first train the adapter before unfreezing the encoder and decoder, keeps gradients smooth and training more stable, preventing this degradation. We therefore adopt the multi-stage recipe for all scaled experiments. We also consider parameter-efficient variants that keep the decoder fully frozen (Ge et al., 2023; Cheng et al., 2024; Yen et al., 2024) or use LoRA (Liao et al., 2025; Feldman and Artzi, 2025), but we find them to substantially underperform compared to full-parameter training in early experiments; we therefore focus on end-to-end training. With compression disabled, the trained decoders show minimal drift from the original model (Appendix D.3). Details of the continual pre-training and SFT data mixtures are provided in Table 1 and Table 2, respectively. All other training details, including token budgets, learning rates, and sequence lengths, are provided in Table 5.

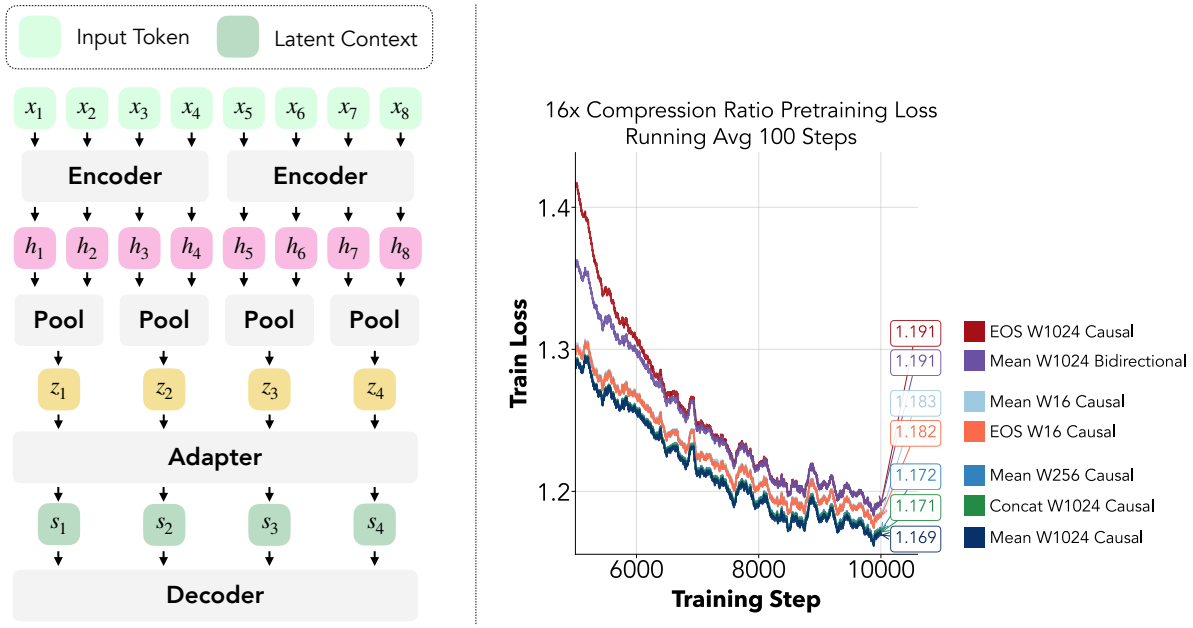


Figure 3 A from-scratch pre-training sweep identifies the best encoder-decoder compressor architecture. Left: We explore how to best train an encoder LM to compress raw text into latent vectors and a decoder LM that uses the vectors as latent context. **Right:** We sweep many variants of encoder-decoder compressors end-to-end with 38B training tokens at a compression ratio of 16× and compare pre-training loss to find the best architecture. Since loss is computed only on the uncompressed tokens, the pre-training loss is substantially lower than a standard next-token prediction loss over all tokens. Extended experiments are in Appendix C.

5 Architectural Design Space of Latent Encoders

We conduct a controlled architecture search over encoder-decoder context compressors by pre-training LCLM variants from scratch. Our cleanroom setting lets us isolate the effect of architectural choices without confounding from pre-trained initialization. This is particularly important for choices such as the encoder attention mask and pooling operators: for example, the pre-trained `Qwen3-Embedding-0.6B` encoder uses causal attention and EOS pooling, which could bias the architectural search. Specifically, we use `Qwen3-0.6B` as the base architecture for both the encoder and decoder, randomly initializing all weights with the same seed, and training all components end-to-end with no frozen parameters. Each variant is trained on the pre-training mixture described in Section 4 for 38B tokens at a compression ratio $N = 16$. This comprehensive search covers nearly all major architectural choices explored in prior soft-token compression work, along with additional underexplored choices. Prior soft-token compressors can be decomposed into a small set of design choices: the pooling operator, encoding granularity, encoder attention mask, adapter, encoder initialization, and decoder adaptation. Each prior method is a point in this design space (Table 4 in Appendix C.1), so sweeping each axis under matched data, scale, and decoder compares the choices behind prior methods without the confounding differences between their original setups. We describe each design axis and its empirical effect below.

Pooling operator. Prior work on context compression has primarily focused on two pooling strategies. Token-based pooling appends or prepends designated tokens, such as `EOS` or `CLS`, and uses their final hidden states as compressed representations (Chevalier et al., 2023; Ge et al., 2023). Mean pooling averages encoder hidden states over each block of N input tokens (Tang et al., 2025; Feldman and Artzi, 2025), a design that is common in vision and multimodal models (Gemma Team et al., 2026). We additionally consider concatenation, a widely used alternative in vision-style architectures but less explored for context compression (Bai et al., 2025; Tong et al., 2024). Concatenation preserves the N sequential encoder hidden states by stacking them into a single wide vector, after which the adapter projects the resulting representation back to the LLM hidden dimension. The formal definitions

of pooling operators can be found in Appendix C.2.

Figure 3 compares the three pooling operators in our pre-training sweep. Mean pooling matches or improves pre-training loss over token-based pooling, and performs on par with concatenation. This agrees with prior evidence that mean pooling can outperform token-based pooling for compression (Feldman and Artzi, 2025), and with findings that mean pooling can be surprisingly effective despite being a lossy operation in text embedding models (Hara et al., 2026). It is also consistent with related observations in vision architectures (Chu et al., 2021). While concatenation can, in principle, preserve distinct representations that mean loses, and has been found to outperform mean pooling in other representation-learning settings (Chen et al., 2025), we do not observe this advantage in our experiments. We compare these two operators at scale in Figure 9 and Appendix F.6.1, where concatenation outperforms mean pooling at low compression ratios and mean pooling outperforms concatenation at high compression ratios.

Encoding granularity. Ideally, one would set the encoder window size to the full context length, allowing the encoder to jointly contextualize all input tokens, analogous to ViT-style encoders (Dosovitskiy et al., 2020) that process the entire image at once. For language inputs, however, the context length can be substantially larger, making full-context encoding prohibitive in both memory and time.

To study this trade-off, we vary the encoder window size over $W \in \{N, 256, 1024\}$ for mean pooling in Figure 3. We find that increasing W from N (16) (Lin et al., 2025; Liao et al., 2025) to 256 leads to a large improvement, while increasing W further from 256 to 1024 yields an additional smaller gain. Since $W = 1024$ does not introduce significant memory or runtime overhead in our setting, we adopt $W = 1024$ as the default. This setting enables the encoder to attend over a larger local context and produce richer compressed representations. One limitation of using $W < T$ is that the input must be chunked into multiple windows, so information crossing a boundary is split across two encoder forward passes. To address this, we also design a boundary-overlap variant that allows each window to see neighboring tokens during encoding. However, this approach does not consistently improve pre-training loss while increasing encoder compute, so we use no overlap by default and defer details to Appendix C.4.

Encoder attention mask. Unlike standard decoder language modeling, where causal masking is required to match autoregressive inference, context compression is applied to the prompt before decoding. The encoder can therefore use either causal or bidirectional attention, as in prefix-LMs (Raffel et al., 2020). Prior compressors use both: some encode with a causal LLM (Ge et al., 2023; Tang et al., 2025), while others remove the causal mask (Feldman and Artzi, 2025). In Figures 3 and 10, we compare these two masking choices and find that causal masking consistently achieves lower pre-training loss.

Adapter design. The adapter projects encoder latents from dimension d_{enc} to the decoder embedding dimension d_{dec}. Prior encoder-decoder compression work often assumes matching encoder and decoder hidden sizes, but a wide range of encoder-decoder models pair smaller encoders with larger decoders, making the adapter a non-trivial design choice. Prior work uses either a small MLP projector (Cheng et al., 2024; Lin et al., 2025) or attention-based layers (Tang et al., 2025). We compare a lightweight MLP projection against an attention-based adapter that adds a single self-attention layer over the latent sequence before the MLP projection. In contrast to prior findings (Tang et al., 2025), we observe in Figure 10 that the MLP-only adapter achieves lower pre-training loss while using less compute, and we adopt it as the default. Formal definitions are provided in Appendix C.3.

Optimal architecture at scale. Overall, our from-scratch pre-training experiments in Figure 3 serve as a lightweight architecture sweep, allowing us to narrow the search space to a smaller set of promising design choices. We then scale these candidates using the full training pipeline, both to identify the strongest architecture at scale and to validate that the small-scale pre-training findings remain predictive under large-scale training. For our main experiments, we use `Qwen3-Embedding-0.6B` as the encoder with `Qwen3-4B-Instruct-2507` as the decoder. Training at scale confirms the choices from the small-scale sweep: both the continual pre-training losses (Figures 9 to 11) and the downstream ablations (Appendix F.6) again favor an encoder window $W = 1024$ and causal masking. At scale, the attention-based adapter is competitive with the MLP adapter; we keep the MLP adapter, which adds less computation. The choice of pooling operator depends on the compression ratio: mean pooling has a small edge over concatenation at 16× compression, while concatenation slightly outperforms mean pooling at 4×. Since we primarily focus on 16× compression in our analysis, we adopt mean pooling as the default in all subsequent experiments. Full results and per-task breakdowns are reported in Appendix F.6.

We also study alternative model choices, including whether to initialize the encoder from an embedding model (Liao et al., 2025; Cheng et al., 2024) or a language model (Ge et al., 2023; Tang et al., 2025; Feldman and Artzi, 2025), in Appendix D.1. In Appendix D.4, we analyze the performance of scaling the sizes of the encoder and decoder.

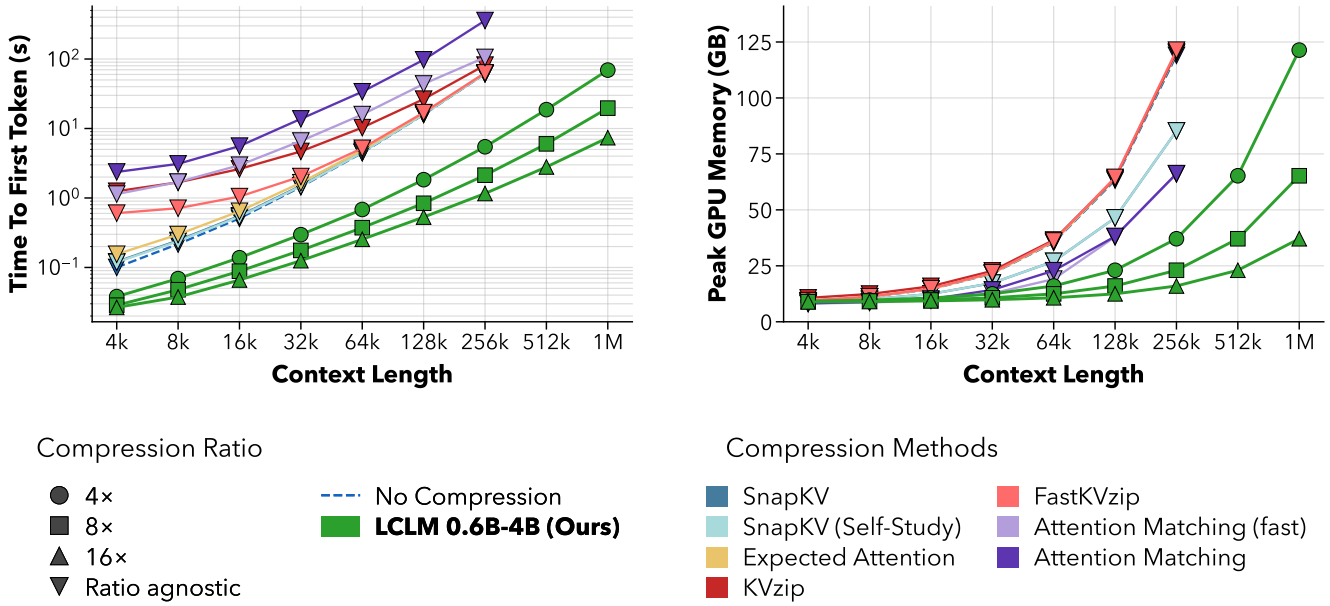


Figure 4 Latent Context Language Models have lower TTFT and peak GPU memory as context length increases. We plot time to first token (left) and peak GPU memory (right) for a single sample. We note the baselines use the same time and memory regardless of compression ratio; hence, there is only one line per baseline method. We record all measurements on the same H200 GPU, and truncate lines when the method cannot be performed within the 141GB of memory. For peak GPU memory, our encoder's activations stay small, so memory grows only with the decoder prefill over the compressed latent sequence.

6 Results: Improving Latency / Memory-Performance Tradeoffs

We predominantly focus on long-context benchmarks: we evaluate on RULER (Hsieh et al., 2024), LongBench (Bai et al., 2024), and LongHealth (Adams et al., 2025). For the long-context evaluation suite, we maintain instructions as uncompressed tokens, while we compress long-context segments. We report which parts of each benchmark are compressed and which are left uncompressed in Table 8.

We measure two efficiency axes: peak memory during compression, and compression time. Peak memory includes both the encoder and decoder for a single sample at each context length. For every method, we report the fastest of 20 timed runs after 30 warmup runs on the same H200 GPU. We define compression time as time-to-first-token (TTFT), i.e., the time required for a method to reach the point at which it can generate the first token. We adapt the KVPress (Devoto et al., 2025) timing code[1] to support our method. Because the KV cache compression baselines are implemented in Hugging Face Transformers (Wolf et al., 2019), we benchmark every method in their comparisons with it, including our decoder; Appendix F.3 details the implementation of every timing setup.

Since LCLMs process input tokens in encoder-window chunks, we can parallelize compression by batching encoder forward passes before passing the resulting soft tokens to the decoder. For all figures in this paper, we use an encoder batch size of 32; with encoder window size $W = 1024$, this corresponds to processing 32,768 input tokens per batched encoder pass. We find this batch size provides a good trade-off between compression speed and encoder memory usage, improving GPU utilization without exceeding memory limits for ultra-long contexts. Moreover, because our method preserves the standard KV cache structure, prefill and generation can be further accelerated with standard inference frameworks such as vLLM (Kwon et al., 2023) and SGLang (Zheng et al., 2024), as well as efficient attention implementations (Katharopoulos et al., 2020; Kimi Team et al., 2025; Liu et al., 2024a; Xu et al., 2026). We use this in Figure 5: that comparison involves only LCLMs and uncompressed models, whose decoders all run in vLLM, so we serve every decoder there with vLLM and add our encoder's time for LCLMs. The Hugging Face measurements in Figures 1 and 4 are therefore conservative relative to what optimized serving implementations could achieve.

[1] https://github.com/NVIDIA/kvpress/blob/main/notebooks/speed_and_memory.ipynb

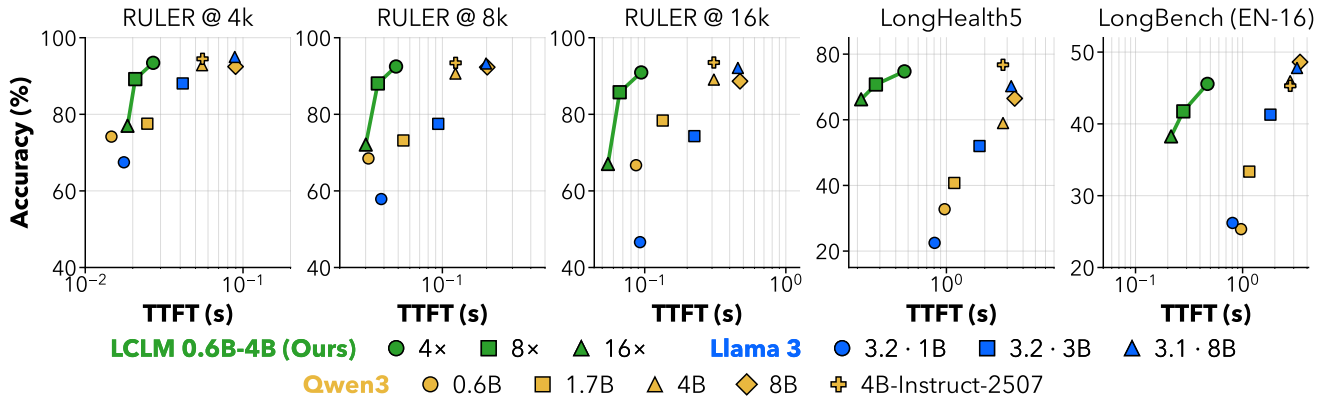


Figure 5 LCLMs are faster and more accurate than smaller uncompressed language models. Accuracy vs. TTFT for LCLMs and uncompressed Qwen3 and Llama 3 models from 0.6B to 8B parameters. All models are optimized for inference, with every decoder served by vLLM on a single H200, timed at fixed context lengths (4k, 8k, and 16k tokens for RULER; 64k for LongHealth5 and LongBench), taking the fastest of 20 trials after 30 warmup runs; LCLM TTFT includes encoding.

6.1 Baselines

We consider a range of KV cache compression baselines, including Expected Attention (Devoto et al., 2025), Attention Matching (Zweiger et al., 2026), KVzip (Kim et al., 2025), FastKVzip (Kim et al., 2026), and SnapKV (Li et al., 2024c). Each method has different native support for query-aware, query-agnostic, and self-study compression configurations, which we describe in Appendix F.2 together with the difference between algorithmic and systems-realizable cache reduction.

We use the KVPress (Devoto et al., 2025) implementation for SnapKV, KVzip, FastKVzip, and Expected Attention. We use the official implementation for Attention Matching.

6.2 A New Pareto Frontier

We train LCLMs with compression ratios of 16×, 8×, and 4×. In Figure 1, LCLMs establish a new Pareto frontier in compression time and accuracy over KV cache compression baselines on RULER, LongBench, and LongHealth. KV cache compression methods first materialize the full KV cache and then evict, mask, or compact it, so their compression time barely depends on the target ratio. In contrast, LCLMs reduce the sequence length before decoder prefill, so higher compression ratios directly reduce decoder-side computation and memory.

Comparison with smaller uncompressed models. A simple alternative to compression is to run a smaller model on the full context. In Figure 5, we compare LCLMs with uncompressed Qwen3 (Yang et al., 2025) and Llama 3 (Grattafiori et al., 2024) models from 0.6B to 8B parameters, with every model optimized for inference. At 64k tokens, the 16× LCLM reaches the first token faster than the smallest uncompressed models (0.21 s, compared with 0.81 s for Llama 3.2 1B and 0.97 s for Qwen3-0.6B) while being far more accurate on LongHealth5 (66.25% vs. 22.50% and 32.75%). The 4× LCLM exceeds the uncompressed `Qwen3-4B-Instruct-2507` on LongBench with a 6× lower TTFT.

Time and memory scaling with context length. In Figure 4, we visualize compression time and peak GPU memory from 4K to 1M tokens. LCLMs achieve the fastest compression time at every context length and substantially lower peak memory at long contexts; all baselines run out of memory at 512K tokens, except Attention Matching, which fails at 512K due to numerical instability in its linear solver and runs out of memory at 1M. Because the encoder is small and processes at most 32K input tokens per batched forward pass, peak memory is driven by the decoder prefill over the compressed latent sequence: the 16× model grows from 9 GB at 4K tokens to only 23 GB at 512K, and at 1M tokens the 16×, 8×, and 4× models peak at 37, 65, and 121 GB.

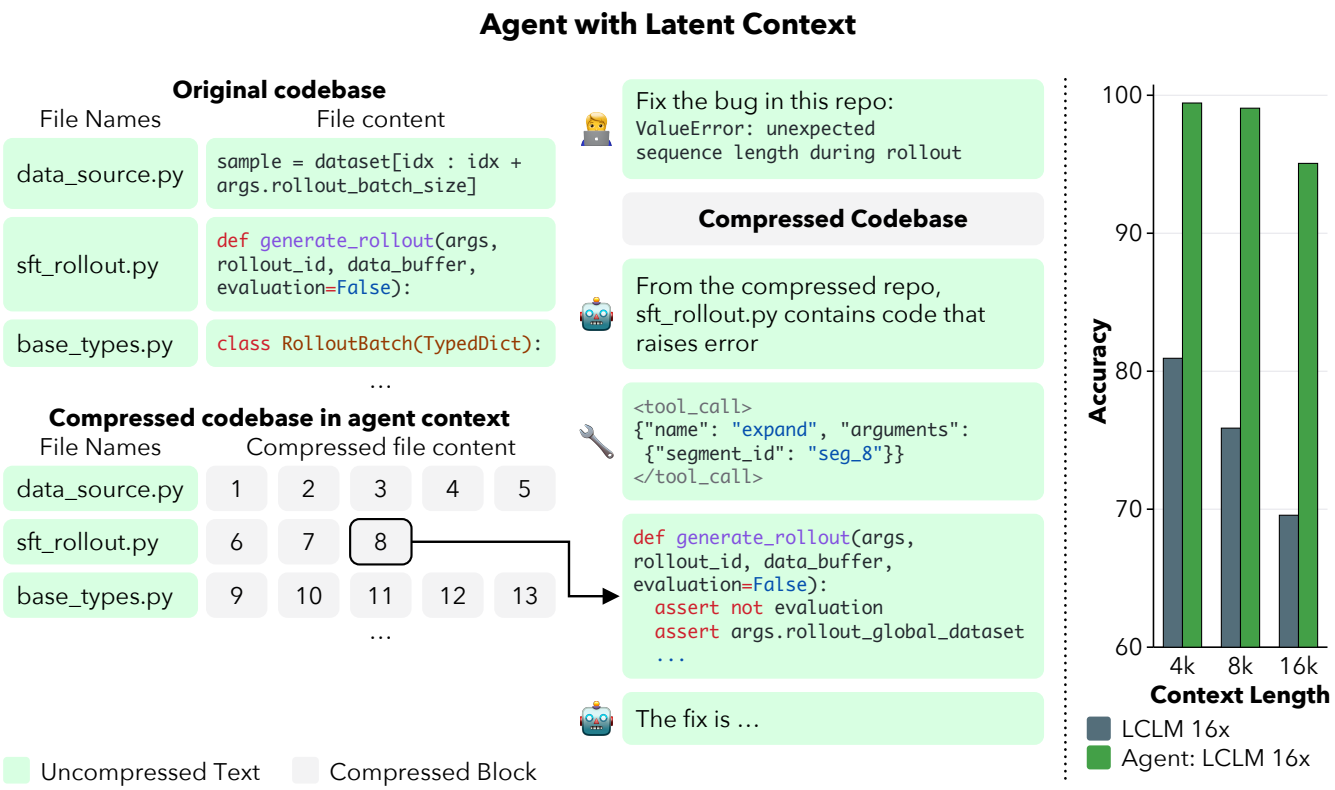


Figure 6 LCLMs can use tools to retrieve compressed context and improve retrieval accuracy. Left: We equip LCLMs with a tool that retrieves compressed context and inserts it into their working context uncompressed. **Right:** RULER NIAH scores at 16× compression, with and without expansion, averaged over eight tasks with 50 matched questions per task and context length. We cap the agent at 16 expansion calls at 4k and 8k and 32 at 16k; these caps were chosen on three of the eight tasks and kept fixed for the rest.

7 New Capabilities from Compression

Because the decoder reads latent tokens in place of the original text, LCLMs enable uses beyond compressing a fixed input. We study two: an agent that reads a compressed context and expands segments on demand (Section 7.1), and online latent compaction of the model's own generated reasoning (Section 7.2).

7.1 Agent Workflow With Latent Context

Agentic systems are commonly used to handle large amounts of information (Zhang et al., 2025a). However, agents usually rely on lexical or semantic search to locate the information needed. These information retrieval methods can fail when the search keyword is not obvious before reading through the information. For example, in a large codebase, a bug reported in the "dashboard login flow" may actually originate in an indirectly called entitlement module that never mentions "dashboard" or "login". Ideally, an agent should read all information at once before acting, but context limits of existing models prevent this from happening.

LCLMs' more efficient long-context capacity makes it possible to place, for example, an entire code repository in an agent's context. We further enhance the agent's capability to locate precise information by allowing the LCLM agent to expand compressed segments into raw text on demand. We segment the input into fixed-size chunks of 1,024 encoder tokens, compress each chunk, and assign it a segment identifier. The model receives the entire compressed sequence in a single prompt, together with one `expand` tool. At each turn, the model can issue a native tool call to `expand` with a segment identifier (Figure 6), and the tool returns the original text of that segment.

We use the needle-in-a-haystack tasks from RULER as a testbed for evaluating this agentic retrieval mechanism with latent context. As shown in Figure 6, our agent substantially improves performance over the same LCLM with 16× compressed context and no expansion. These results suggest that compressed latent context can provide broad

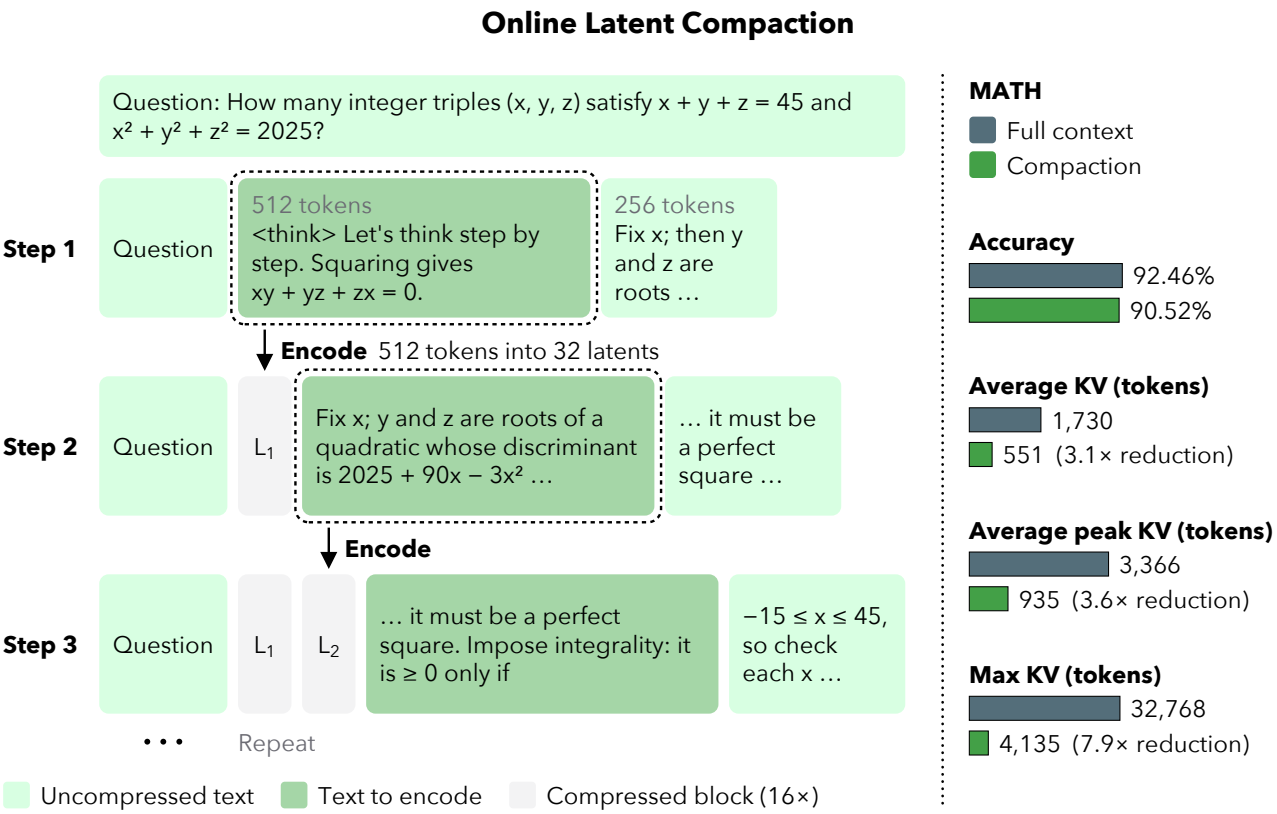


Figure 7 **Online latent compaction. Left:** During generation, every 512 uncompressed tokens are encoded into 32 latent tokens (16× compression), while the prompt and the latest 256 tokens stay uncompressed. We iteratively compress and generate from the compacted context and the latest raw tokens. **Right:** Accuracy and decoder context length on the full MATH test set. Average KV and peak KV are computed per problem and averaged over problems; max KV is the largest per-problem peak, where the full-context limit is 32,768 tokens.

corpus-level visibility, while tool-based expansion supplies the exact fine-grained information needed. The marked success of this initial agentic attempt suggests the promise of adaptive expansion. LCLMs can skim a large quantity of context globally before deciding which relevant subset to zoom in on and read more carefully. Future work may explore how to choose which context to expand, or learn this behavior end-to-end.

7.2 Online Latent Compaction

For long-horizon tasks, long CoT and agent trajectories can come to dominate the context and eventually fill the context window, requiring context compaction (Anthropic, 2025; OpenAI, 2025). Here, we investigate whether we can use LCLMs to compact context into the latent space iteratively. Because the LCLM encoder can compress any span of text, we can also apply it to the model's own output during generation, which we call online latent compaction (Figure 7). After the first 768 generated tokens, we encode the oldest 512 generated tokens into latent memory tokens at a 16× compression ratio, resulting in 32 latent tokens, and keep the most recent 256 tokens uncompressed. Generation then continues from the compacted context and latest raw tokens, and the procedure repeats after every further 512 generated tokens. Earlier latent blocks are reused thereafter.

We evaluate on all 5,000 problems of the MATH test set with the 16× LCLM, comparing generation with and without compaction under the same prompts, 32,768-token context limit, and scoring. Online latent compaction is nearly lossless, reaching 90.52% accuracy compared with 92.46% for full-context generation, while being highly efficient: it reduces the average KV length by 3.1×, the average per-problem peak by 3.6×, and the maximum peak by 7.9×: the longest full-context trajectories reach the 32,768-token limit, whereas no compacted trajectory exceeds 4,135 tokens. Lengths count active decoder positions, including latents, not allocated memory.

8 Conclusions

Through a controlled architecture search and end-to-end training at scale, our family of Latent Context Language Models shows that encoder-decoder compression can preserve strong in-context capabilities while compressing faster and using less memory than KV cache compression. The same models support agents that expand compressed context on demand and online compaction of the model's own reasoning, making them a practical building block for efficient long-context systems.

Latent Context Language Models open a broad design space for future work and are naturally compatible with agentic frameworks such as Recursive Language Models (Zhang et al., 2025a). More generally, LCLMs provide a promising substrate for long-horizon agents with large, persistent working memories by dramatically reducing the size of the inputs at scale. Future iterations of LCLMs could further improve the quality-efficiency trade-off by compressing inputs at multiple granularities and dynamically allocating capacity based on information density or input perplexity. Such adaptive compression could allow models to preserve fine-grained details where needed while maintaining a compact global context, reducing reliance on explicit expansion tools. Another promising direction is to extend online latent compaction beyond chain-of-thought to tool observations and an agent's accumulated working history, which can grow to dominate the context budget in long-horizon tasks.

Acknowledgments

MG and ZC were supported by IBM, the Google Research Award, and the NVIDIA Academic Grant Program for Researchers. SM and TG are supported by DOE Office of Science's ASCR AI for Science initiative, the NSF TRAILS Institute (2229885), and Coefficient Giving. NK is supported by the NSF Graduate Research Fellowship.

Computing resources for this project were supported in part by the Swiss National Supercomputing Center (CSCS). Prepared in collaboration with LLNL under Contract DE-AC52-07NA27344 and supported by the LLNL-LDRD Program under Project No. 24-ERD-010 (LLNL-CONF-2019739). This manuscript has been authored by Lawrence Livermore National Security, LLC under Contract No. DE-AC52-07NA27344 with the U.S. Department of Energy. The United States Government retains, and the publisher, by accepting the article for publication, acknowledges that the United States Government retains a non-exclusive, paid-up, irrevocable, world-wide license to publish or reproduce the published form of this manuscript, or allow others to do so, for United States Government purposes. This material is based upon work supported by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research, through solicitation DE-FOA-0003264, "Advancements in Artificial Intelligence for Science," under Award Number DE-SC0025598.

Special thanks to Shengbang Tong for extensive discussions on architecture design and encoder-decoder model training. We would like to thank Preston Zhou and Eitan Borgnia from Relace AI for their early discussions and support for the project. We also thank Vadim Bereznyuk, John Kirchenbauer, and Neel Jain for providing feedback on the draft.

References

Lisa Adams, Felix Busch, Tianyu Han, Jean-Baptiste Excoffier, Matthieu Ortala, Alexander Löser, Hugo JWL Aerts, Jakob Nikolas Kather, Daniel Truhn, and Keno Bressem. Longhealth: A question answering benchmark with long clinical documents. *Journal of Healthcare Informatics Research*, 9(3):280–296, 2025. 2, 8

Sandhini Agarwal, Lama Ahmad, Jason Ai, Sam Altman, Andy Applebaum, Edwin Arbus, Rahul K Arora, Yu Bai, Bowen Baker, Haiming Bao, et al. gpt-oss-120b & gpt-oss-20b model card. *arXiv preprint arXiv:2508.10925*, 2025. 21

Tushar Aggarwal, Swayam Singh, Abhijeet Awasthi, Aditya Kanade, and Nagarajan Natarajan. Nextcoder: Robust adaptation of code LMs to diverse code edits. In *Forty-second International Conference on Machine Learning*, 2025. URL https://openreview.net/forum?id=3B6fF1PxYD. 22

Chenxin An, Jun Zhang, Ming Zhong, Lei Li, Shansan Gong, Yao Luo, Jingjing Xu, and Lingpeng Kong. Why does the effective context length of LLMs fall short? In *The Thirteenth International Conference on Learning Representations*, 2025. URL https://openreview.net/forum?id=eoln5WgrPx. 1

Anthropic. Claude Code, 2025. URL https://github.com/anthropics/claude-code. 11, 20

Shuai Bai, Yuxuan Cai, Ruizhe Chen, Keqin Chen, Xionghui Chen, Zesen Cheng, Lianghao Deng, Wei Ding, Chang Gao, Chunjiang Ge, et al. Qwen3-VL technical report. *arXiv preprint arXiv:2511.21631*, 2025. 6, 25

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. LongBench: A bilingual, multitask benchmark for long context understanding. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3119–3137, Bangkok, Thailand, August 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.acl-long.172. URL https://aclanthology.org/2024.acl-long.172. 2, 8

Aaron Blakeman, Aaron Grattafiori, Aarti Basant, Abhibha Gupta, Abhinav Khattar, Adi Renduchintala, Aditya Vavre, Akanksha Shukla, Akhiad Bercovich, Aleksander Ficek, et al. Nemotron 3 nano: Open, efficient mixture-of-experts hybrid mamba-transformer model for agentic reasoning. *arXiv preprint arXiv:2512.20848*, 2025. 5, 21, 22

Baiyu Chen, Wilson Wongso, Zechen Li, Yonchanok Khaokaew, Hao Xue, and Flora Salim. Comodo: Cross-modal video-to-imu distillation for efficient egocentric human activity recognition. *arXiv preprint arXiv:2503.07259*, 2025. 7

Longze Chen. Awesome-kv-cache-compression. GitHub repository, 2023. URL https://github.com/October2001/Awesome-KV-Cache-Compression. 20

Xin Cheng, Xun Wang, Xingxing Zhang, Tao Ge, Si-Qing Chen, Furu Wei, Huishuai Zhang, and Dongyan Zhao. xrag: Extreme context compression for retrieval-augmented generation with one token. *Advances in Neural Information Processing Systems*, 37: 109487–109516, 2024. 3, 4, 5, 7, 24

Alexis Chevalier, Alexander Wettig, Anirudh Ajith, and Danqi Chen. Adapting language models to compress contexts. *arXiv preprint arXiv:2305.14788*, 2023. 2, 4, 6, 20, 24, 27

Xiangxiang Chu, Zhi Tian, Bo Zhang, Xinlong Wang, and Chunhua Shen. Conditional positional encodings for vision transformers. *arXiv preprint arXiv:2102.10882*, 2021. 7

Yu-Neng Chuang, Tianwei Xing, Chia-Yuan Chang, Zirui Liu, Xun Chen, and Xia Hu. Learning to compress prompt in natural language formats. *arXiv preprint arXiv:2402.18700*, 2024. 20

Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, et al. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*, 2021. 31

Domenico Cotroneo, Giuseppe De Rosa, and Pietro Liguori. Pyresbugs: A dataset of residual python bugs for natural language-driven fault injection. In *2025 IEEE/ACM Second International Conference on AI Foundation Models and Software Engineering (Forge)*, pages 146–150, 2025. doi: 10.1109/Forge66646.2025.00024. 22

Yuhong Dai, Jianxun Lian, Yitian Huang, Wei Zhang, Mingyang Zhou, Mingqi Wu, Xing Xie, and Hao Liao. Pretraining context compressor for large language models with embedding-based memory. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 28715–28732, 2025. 3, 20

Zihang Dai, Zhilin Yang, Yiming Yang, Jaime G Carbonell, Quoc Le, and Ruslan Salakhutdinov. Transformer-xl: Attentive language models beyond a fixed-length context. In *Proceedings of the 57th annual meeting of the association for computational linguistics*, pages 2978–2988, 2019. 20

Tri Dao. Flashattention-2: Faster attention with better parallelism and work partitioning. *arXiv preprint arXiv:2307.08691*, 2023. 29, 32

Alessio Devoto, Maximilian Jeblick, and Simon Jégou. Expected attention: Kv cache compression by estimating attention from future queries distribution. *arXiv preprint arXiv:2510.00636*, 2025. 1, 3, 8, 9, 20

Juechu Dong, Boyuan Feng, Driss Guessous, Yanbo Liang, and Horace He. Flex attention: A programming model for generating optimized attention kernels. *arXiv preprint arXiv:2412.05496*, 2(3):4, 2024. 29

Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 7

Sabri Eyuboglu, Ryan Ehrlich, Simran Arora, Neel Guha, Dylan Zinsley, Emily Liu, Will Tennien, Atri Rudra, James Zou, Azalia Mirhoseini, et al. Cartridges: Lightweight and general-purpose long context representations via self-study. *arXiv preprint arXiv:2506.06266*, 2025. 1, 3, 20, 32

Yair Feldman and Yoav Artzi. Simple context compression: Mean-pooling and multi-ratio training. *arXiv preprint arXiv:2510.20797*, 2025. 3, 4, 5, 6, 7, 24, 27

Tianyu Gao, Alexander Wettig, Howard Yen, and Danqi Chen. How to train long-context language models (effectively). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 7376–7399, 2025. 5

Tao Ge, Jing Hu, Lei Wang, Xun Wang, Si-Qing Chen, and Furu Wei. In-context autoencoder for context compression in a large language model. *arXiv preprint arXiv:2307.06945*, 2023. 2, 4, 5, 6, 7, 20, 24

Gemma Team, Sherif El Abd, Vaibhav Aggarwal, Robin Algayres, Alek Andreev, Olivier Bachem, Ian Ballantyne, Cormac Brick, Victor Cărbune, Michelle Casbon, et al. Gemma 4 technical report. *arXiv preprint arXiv:2607.02770*, 2026. 6

Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, et al. The Llama 3 herd of models. *arXiv preprint arXiv:2407.21783*, 2024. 9

Albert Gu and Tri Dao. Mamba: Linear-time sequence modeling with selective state spaces. In *First conference on language modeling*, 2024. 21

Albert Gu, Karan Goel, and Christopher Ré. Efficiently modeling long sequences with structured state spaces. *arXiv preprint arXiv:2111.00396*, 2021. 21

Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Peiyi Wang, Qihao Zhu, Runxin Xu, Ruoyu Zhang, Shirong Ma, Xiao Bi, et al. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*, 2025. 21

Tomomasa Hara, Hiroto Kurita, Masaaki Imaizumi, Kentaro Inui, and Sho Yokoi. Why mean pooling works: Quantifying second-order collapse in text embeddings. *arXiv preprint arXiv:2604.27398*, 2026. 7

Linda He, Jue Wang, Maurice Weber, Shang Zhu, Ben Athiwaratkun, and Ce Zhang. Scaling instruction-tuned llms to million-token contexts via hierarchical synthetic data generation. *arXiv preprint arXiv:2504.12637*, 2025. 22

Dan Hendrycks and Kevin Gimpel. Gaussian error linear units (gelus). *arXiv preprint arXiv:1606.08415*, 2016. 26

Coleman Hooper, Sehoon Kim, Hiva Mohammadzadeh, Michael W Mahoney, Yakun S Shao, Kurt Keutzer, and Amir Gholami. Kvquant: Towards 10 million context length llm inference with kv cache quantization. *Advances in Neural Information Processing Systems*, 37:1270–1303, 2024. 1

Cheng-Ping Hsieh, Simeng Sun, Samuel Kriman, Shantanu Acharya, Dima Rekesh, Fei Jia, Yang Zhang, and Boris Ginsburg. Ruler: What's the real context size of your long-context language models? *arXiv preprint arXiv:2404.06654*, 2024. 2, 3, 8

Huiqiang Jiang, Qianhui Wu, Chin-Yew Lin, Yuqing Yang, and Lili Qiu. Llmlingua: Compressing prompts for accelerated inference of large language models. *arXiv preprint arXiv:2310.05736*, 2023. 20

Huiqiang Jiang, Qianhui Wu, Xufang Luo, Dongsheng Li, Chin-Yew Lin, Yuqing Yang, and Lili Qiu. Longllmlingua: Accelerating and enhancing llms in long context scenarios via prompt compression. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1658–1677, 2024. 20

Qiao Jin, Bhuwan Dhingra, Zhengping Liu, William Cohen, and Xinghua Lu. Pubmedqa: A dataset for biomedical research question answering. In *Proceedings of the 2019 conference on empirical methods in natural language processing and the 9th international joint conference on natural language processing (EMNLP-IJCNLP)*, pages 2567–2577, 2019. 22

Mandar Joshi, Eunsol Choi, Daniel S Weld, and Luke Zettlemoyer. Triviaqa: A large scale distantly supervised challenge dataset for reading comprehension. *arXiv preprint arXiv:1705.03551*, 2017. 22

Angelos Katharopoulos, Apoorv Vyas, Nikolaos Pappas, and François Fleuret. Transformers are rnns: Fast autoregressive transformers with linear attention. In *International conference on machine learning*, pages 5156–5165. PMLR, 2020. 8, 21

Jang-Hyun Kim, Jinuk Kim, Sangwoo Kwon, Jae W Lee, Sangdoo Yun, and Hyun Oh Song. Kvzip: Query-agnostic kv cache compression with context reconstruction. *arXiv preprint arXiv:2505.23416*, 2025. 1, 3, 9, 20, 32

Jang-Hyun Kim, Dongyoon Han, and Sangdoo Yun. Fast kvzip: Efficient and accurate llm inference with gated kv eviction. *arXiv preprint arXiv:2601.17668*, 2026. 9, 32

Kimi Team, Yu Zhang, Zongyu Lin, Xingcheng Yao, Jiaxi Hu, Fanqing Meng, Chengyin Liu, Xin Men, Songlin Yang, Zhiyuan Li, et al. Kimi linear: An expressive, efficient attention architecture. *arXiv preprint arXiv:2510.26692*, 2025. 8, 21

Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph Gonzalez, Hao Zhang, and Ion Stoica. Efficient memory management for large language model serving with pagedattention. In *Proceedings of the 29th symposium on operating systems principles*, pages 611–626, 2023. 3, 8, 32

Nathan Lambert, Jacob Morrison, Valentina Pyatkin, Shengyi Huang, Hamish Ivison, Faeze Brahman, Lester James V. Miranda, Alisa Liu, Nouha Dziri, Shane Lyu, Yuling Gu, Saumya Malik, Victoria Graf, Jena D. Hwang, Jiangjiang Yang, Ronan Le Bras, Oyvind Tafjord, Chris Wilhelm, Luca Soldaini, Noah A. Smith, Yizhong Wang, Pradeep Dasigi, and Hannaneh Hajishirzi. Tülu 3: Pushing frontiers in open language model post-training, 2024. 22

Adrian Łańcucki, Konrad Staniszewski, Piotr Nawrot, and Edoardo Maria Ponti. Inference-time hyper-scaling with KV cache compression. *Advances in Neural Information Processing Systems*, 38:9365–9397, 2026. 3

Hongyu Li, Liang Ding, Meng Fang, and Dacheng Tao. Revisiting catastrophic forgetting in large language model tuning. In *Findings of the association for computational linguistics: EMNLP 2024*, pages 4297–4308, 2024a. 28

Xiang Lisa Li and Percy Liang. Prefix-tuning: Optimizing continuous prompts for generation. *arXiv preprint arXiv:2101.00190*, 2021. 20

Yucheng Li, Bo Dong, Frank Guerin, and Chenghua Lin. Compressing context to enhance inference efficiency of large language models. In *Proceedings of the 2023 conference on empirical methods in natural language processing*, pages 6342–6353, 2023. 20

Yucheng Li, Huiqiang Jiang, Qianhui Wu, Xufang Luo, Surin Ahn, Chengruidong Zhang, Amir H Abdi, Dongsheng Li, Jianfeng Gao, Yuqing Yang, et al. Scbench: A kv cache-centric analysis of long-context methods. *arXiv preprint arXiv:2412.10319*, 2024b. 1, 3, 20

Yuhong Li, Yingbing Huang, Bowen Yang, Bharat Venkitesh, Acyr Locatelli, Hanchen Ye, Tianle Cai, Patrick Lewis, and Deming Chen. Snapkv: Llm knows what you are looking for before generation. *Advances in Neural Information Processing Systems*, 37: 22947–22970, 2024c. 1, 3, 9, 20

Zongqian Li, Yixuan Su, and Nigel Collier. 500xcompressor: Generalized prompt compression for large language models. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 25081–25091, 2025. 20

Zihan Liao, Jun Wang, Hang Yu, Lingxiao Wei, Jianguo Li, and Wei Zhang. E2llm: Encoder elongated large language models for long-context understanding and reasoning. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 19212–19241, 2025. 2, 4, 5, 7, 20, 24, 27

Xiaoqiang Lin, Aritra Ghosh, Bryan Kian Hsiang Low, Anshumali Shrivastava, and Vijai Mohan. Refrag: Rethinking rag based decoding. *arXiv preprint arXiv:2509.01092*, 2025. 2, 3, 4, 7, 24, 27

Aixin Liu, Bei Feng, Bin Wang, Bingxuan Wang, Bo Liu, Chenggang Zhao, Chengqi Dengr, Chong Ruan, Damai Dai, Daya Guo, et al. Deepseek-v2: A strong, economical, and efficient mixture-of-experts language model. *arXiv preprint arXiv:2405.04434*, 2024a. 8, 21

Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. *Advances in neural information processing systems*, 36:34892–34916, 2023a. 5, 26

Nelson F Liu, Kevin Lin, John Hewitt, Ashwin Paranjape, Michele Bevilacqua, Fabio Petroni, and Percy Liang. Lost in the middle: How language models use long contexts. *Transactions of the association for computational linguistics*, 12:157–173, 2024b. 1

Tianyang Liu, Canwen Xu, and Julian McAuley. Repobench: Benchmarking repository-level code auto-completion systems. *arXiv preprint arXiv:2306.03091*, 2023b. 22

Wanlong Liu, Junying Chen, Ke Ji, Li Zhou, Wenyu Chen, and Benyou Wang. Rag-instruct: Boosting llms with diverse retrieval-augmented instructions, 2024c. URL https://arxiv.org/abs/2501.00353. 22

Zihan Liu, Wei Ping, Rajarshi Roy, Peng Xu, Chankyu Lee, Mohammad Shoeybi, and Bryan Catanzaro. Chatqa: Building gpt-4 level conversational qa models. *CoRR*, 2024d. 22

Anton Lozhkov, Raymond Li, Loubna Ben Allal, Federico Cassano, Joel Lamy-Poirier, Nouamane Tazi, Ao Tang, Dmytro Pykhtar, Jiawei Liu, Yuxiang Wei, Tianyang Liu, Max Tian, Denis Kocetkov, Arthur Zucker, Younes Belkada, Zijian Wang, Qian Liu, Dmitry Abulkhanov, Indraneil Paul, Zhuang Li, Wen-Ding Li, Megan Risdal, Jia Li, Jian Zhu, Terry Yue Zhuo, Evgenii Zheltonozhskii, Nii Osae Osae Dade, Wenhao Yu, Lucas Krauß, Naman Jain, Yixuan Su, Xuanli He, Manan Dey, Edoardo Abati, Yekun Chai, Niklas Muennighoff, Xiangru Tang, Muhtasham Oblokulov, Christopher Akiki, Marc Marone, Chenghao Mou, Mayank Mishra, Alex Gu, Binyuan Hui, Tri Dao, Armel Zebaze, Olivier Dehaene, Nicolas Patry, Canwen Xu, Julian McAuley, Han Hu, Torsten Scholak, Sebastien Paquet, Jennifer Robinson, Carolyn Jane Anderson, Nicolas Chapados, Mostofa Patwary, Nima Tajbakhsh, Yacine Jernite, Carlos Muñoz Ferrandis, Lingming Zhang, Sean Hughes, Thomas Wolf, Arjun Guha, Leandro von Werra, and Harm de Vries. Starcoder 2 and the stack v2: The next generation, 2024. 22

Yun Luo, Zhen Yang, Fandong Meng, Yafu Li, Jie Zhou, and Yue Zhang. An empirical study of catastrophic forgetting in large language models during continual fine-tuning. *IEEE Transactions on Audio, Speech and Language Processing*, 2025. 28

Jesse Mu, Xiang Li, and Noah Goodman. Learning to compress prompts with gist tokens. *Advances in Neural Information Processing Systems*, 36:19327–19352, 2023. 20

Niklas Muennighoff, Qian Liu, Armel Zebaze, Qinkai Zheng, Binyuan Hui, Terry Yue Zhuo, Swayam Singh, Xiangru Tang, Leandro von Werra, and Shayne Longpre. Octopack: Instruction tuning code large language models. *arXiv preprint arXiv:2308.07124*, 2023. 22

Dhruv Nathawani, Shuoyang Ding, Vitaly Lavrukhin, Igor Gitman, Somshubra Majumdar, Evelina Bakhturina, Boris Ginsburg, and Jane Polak Scowcroft. Nemotron-Post-Training-Dataset-v2, aug 2025a. URL https://huggingface.co/datasets/nvidia/Nemotron-Post-Training-Dataset-v2. 22

Dhruv Nathawani, Igor Gitman, Somshubra Majumdar, Evelina Bakhturina, Ameya Sunil Mahabaleshwarkar, Jian Zhang, and Jane Polak Scowcroft. Nemotron-Post-Training-Dataset-v1, July 2025b. URL https://huggingface.co/datasets/nvidia/Nemotron-Post-Training-Dataset-v1. 22

Piotr Nawrot, Adrian Łańcucki, Marcin Chochowski, David Tarjan, and Edoardo Ponti. Dynamic memory compression: Retrofitting LLMs for accelerated inference. In Ruslan Salakhutdinov, Zico Kolter, Katherine Heller, Adrian Weller, Nuria Oliver, Jonathan Scarlett, and Felix Berkenkamp, editors, *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pages 37396–37412. PMLR, 21–27 Jul 2024. URL https://proceedings.mlr.press/v235/nawrot24a.html. 3

NVIDIA, Aarti Basant, Abhijit Khairnar, Abhijit Paithankar, Abhinav Khattar, Adithya Renduchintala, Aditya Malte, Akhiad Bercovich, Akshay Hazare, Alejandra Rico, Aleksander Ficek, Alex Kondratenko, Alex Shaposhnikov, Alexander Bukharin, Ali Taghibakhshi, Amelia Barton, Ameya Sunil Mahabaleshwarkar, Amy Shen, Andrew Tao, Ann Guan, Anna Shors, Anubhav Mandarwal, Arham Mehta, Arun Venkatesan, Ashton Sharabiani, Ashwath Aithal, Ashwin Poojary, Ayush Dattagupta, Balaram Buddharaju, Banghua Zhu, Barnaby Simkin, Bilal Kartal, Bita Darvish Rouhani, Bobby Chen, Boris Ginsburg, Brandon Norick, Brian Yu, Bryan Catanzaro, Charles Wang, Charlie Truong, Chetan Mungekar, Chintan Patel, Chris Alexiuk, Christian Munley, Christopher Parisien, Dan Su, Daniel Afrimi, Daniel Korzekwa, Daniel Rohrer, Daria Gitman, David Mosallanezhad, Deepak Narayanan, Dima Rekesh, Dina Yared, Dmytro Pykhtar, Dong Ahn, Duncan Riach, Eileen Long, Elliott Ning, Eric Chung, Erick Galinkin, Evelina Bakhturina, Gargi Prasad, Gerald Shen, Haifeng Qian, Haim Elisha, Harsh Sharma, Hayley Ross, Helen Ngo, Herman Sahota, Hexin Wang, Hoo Chang Shin, Hua Huang, Iain Cunningham, Igor Gitman, Ivan Moshkov, Jaehun Jung, Jan Kautz, Jane Polak Scowcroft, Jared Casper, Jian Zhang, Jiaqi Zeng, Jimmy Zhang, Jinze Xue, Jocelyn Huang, Joey Conway, John Kamalu, Jonathan Cohen, Joseph Jennings, Julien Veron Vialard, Junkeun Yi, Jupinder Parmar, Kari Briski, Katherine Cheung, Katherine Luna, Keith Wyss, Keshav Santhanam, Kezhi Kong, Krzysztof Pawelec, Kumar Anik, Kunlun Li, Kushan Ahmadian, Lawrence McAfee, Laya Sleiman, Leon Derczynski, Luis Vega, Maer Rodrigues de Melo, Makesh Narsimhan Sreedhar, Marcin Chochowski, Mark Cai, Markus Kliegl, Marta Stepniewska-Dziubinska, Matvei Novikov, Mehrzad Samadi, Meredith Price, Meriem Boubdir, Michael Boone, Michael Evans, Michal Bien, Michal Zawalski, Miguel Martinez, Mike Chrzanowski, Mohammad Shoeybi, Mostofa Patwary, Namit Dhameja, Nave Assaf, Negar Habibi, Nidhi Bhatia, Nikki Pope, Nima Tajbakhsh, Nirmal Kumar Juluru, Oleg Rybakov, Oleksii Hrinchuk, Oleksii Kuchaiev, Oluwatobi Olabiyi, Pablo Ribalta, Padmavathy Subramanian, Parth Chadha, Pavlo Molchanov, Peter Dykas, Peter Jin, Piotr Bialecki, Piotr Januszewski, Pradeep Thalasta, Prashant Gaikwad, Prasoon Varshney, Pritam Gundecha, Przemek Tredak, Rabeeh Karimi Mahabadi, Rajen Patel, Ran El-Yaniv, Ranjit Rajan, Ria Cheruvu, Rima Shahbazyan, Ritika Borkar, Ritu Gala, Roger Waleffe, Ruoxi Zhang, Russell J. Hewett, Ryan Prenger, Sahil Jain, Samuel Kriman, Sanjeev Satheesh, Saori Kaji, Sarah Yurick, Saurav Muralidharan, Sean Narenthiran, Seonmyeong Bak, Sepehr Sameni, Seungju Han, Shanmugam Ramasamy, Shaona Ghosh, Sharath Turuvekere Sreenivas, Shelby Thomas, Shizhe Diao, Shreya Gopal, Shrimai Prabhumoye, Shubham Toshniwal, Shuoyang Ding, Siddharth Singh, Siddhartha Jain, Somshubra Majumdar, Soumye Singhal, Stefania Alborghetti, Syeda Nahida Akter, Terry Kong, Tim Moon, Tomasz Hliwiak, Tomer Asida, Tony Wang, Tugrul Konuk, Twinkle Vashishth, Tyler Poon, Udi Karpas, Vahid Noroozi, Venkat Srinivasan, Vijay Korthikanti, Vikram Fugro, Vineeth Kalluru, Vitaly Kurin, Vitaly Lavrukhin, Wasi Uddin Ahmad, Wei Du, Wonmin Byeon, Ximing Lu, Xin Dong, Yashaswi Karnati, Yejin Choi, Yian Zhang, Ying Lin, Yonggan Fu, Yoshi Suhara, Zhen Dong, Zhiyu Li, Zhongbo Zhu, and Zijia Chen. Nvidia nemotron nano 2: An accurate and efficient hybrid mamba-transformer reasoning model, 2025. URL https://arxiv.org/abs/2508.14444. 21

OLMo Team, Allyson Ettinger, Amanda Bertsch, Bailey Kuehl, David Graham, David Heineman, Dirk Groeneveld, Faeze Brahman, Finbarr Timbers, Hamish Ivison, Jacob Morrison, Jake Poznanski, Kyle Lo, Luca Soldaini, Matt Jordan, Mayee Chen, Michael Noukhovitch, Nathan Lambert, Pete Walsh, Pradeep Dasigi, Robert Berry, Saumya Malik, Saurabh Shah, Scott Geng, Shane Arora, Shashank Gupta, Taira Anderson, Teng Xiao, Tyler Murray, Tyler Romero, Victoria Graf, Akari Asai, Akshita Bhagia, Alexander Wettig, Alisa Liu, Aman Rangapur, Chloe Anastasiades, Costa Huang, Dustin Schwenk, Harsh Trivedi, Ian Magnusson, Jaron Lochner, Jiacheng Liu, Lester James V. Miranda, Maarten Sap, Malia Morgan, Michael Schmitz, Michal Guerquin, Michael Wilson, Regan Huff, Ronan Le Bras, Rui Xin, Rulin Shao, Sam Skjonsberg, Shannon Zejiang Shen, Shuyue Stella Li, Tucker Wilde, Valentina Pyatkin, Will Merrill, Yapei Chang, Yuling Gu, Zhiyuan Zeng, Ashish Sabharwal, Luke Zettlemoyer, Pang Wei Koh, Ali Farhadi, Noah A. Smith, and Hannaneh Hajishirzi. Olmo 3, 2025. URL https://arxiv.org/abs/2512.13961. 21, 22

OpenAI. Introducing codex, May 2025. URL https://openai.com/index/introducing-codex/. Accessed: 2026-01-09. 11, 20

Charles Packer, Sarah Wooders, Kevin Lin, Vivian Fang, Shishir G. Patil, Ion Stoica, and Joseph E. Gonzalez. Memgpt: Towards LLMs as operating systems, 2023. URL https://arxiv.org/abs/2310.08560. 20

Guilherme Penedo. Finewiki, 2025. URL https://huggingface.co/datasets/HuggingFaceFW/finewiki. Source: Wikimedia Enterprise Snapshot API (https://api.enterprise.wikimedia.com/v2/snapshots). Text licensed under CC BY-SA 4.0 with attribution to Wikipedia contributors. 22

Bowen Peng, Jeffrey Quesnelle, Honglu Fan, and Enrico Shippole. Yarn: Efficient context window extension of large language models. *arXiv preprint arXiv:2309.00071*, 2023. 21

Hippolyte Pilchen, Edouard Grave, and Patrick Pérez. Arc-encoder: learning compressed text representations for large language models. *arXiv preprint arXiv:2510.20535*, 2025. 4

Jack W Rae, Anna Potapenko, Siddhant M Jayakumar, and Timothy P Lillicrap. Compressive transformers for long-range sequence modelling. *arXiv preprint arXiv:1911.05507*, 2019. 20

Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of machine learning research*, 21(140): 1–67, 2020. 7

Negin Raoof, Richard Zhuang, Marianna Nezhurina, Etash Guha, Atula Tejaswi, Ryan Marten, Charlie F Ruan, Tyler Griggs, Alexander Glenn Shaw, Hritik Bansal, et al. OpenThoughts-Agent: Data recipes for agentic models. *arXiv preprint arXiv:2606.24855*, 2026. 5, 22

Jianlin Su, Murtadha Ahmed, Yu Lu, Shengfeng Pan, Wen Bo, and Yunfeng Liu. Roformer: Enhanced transformer with rotary position embedding. *Neurocomputing*, 568:127063, 2024. 21

Ryan Synk, Monte Hoover, John Kirchenbauer, Neel Jain, Alex Stein, Manli Shu, Josue Melendez Sanchez, Ramani Duraiswami, and Tom Goldstein. Exploiting sparsity for long context inference: Million token contexts on commodity gpus. *arXiv preprint arXiv:2502.06766*, 2025. 20

Sijun Tan, Xiuyu Li, Shishir G Patil, Ziyang Wu, Tianjun Zhang, Kurt Keutzer, Joseph E Gonzalez, and Raluca Ada Popa. Lloco: Learning long contexts offline. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 17605–17621, 2024. 20

Jiwei Tang, Zhicheng Zhang, Shunlong Wu, Jingheng Ye, Lichen Bai, Zitai Wang, Tingwei Lu, Jiaqi Chen, Lin Hai, Hai-Tao Zheng, et al. Gmsa: Enhancing context compression via group merging and layer semantic alignment. *arXiv preprint arXiv:2505.12215*, 2025. 3, 4, 6, 7, 24, 26, 27

Shengbang Tong, Ellis Brown, Penghao Wu, Sanghyun Woo, Manoj Middepogu, Sai C Akula, Jihan Yang, Shusheng Yang, Adithya Iyer, Xichen Pan, et al. Cambrian-1: A fully open, vision-centric exploration of multimodal llms. *Advances in Neural Information Processing Systems*, 37:87310–87356, 2024. 5, 6

Bing Wang, Xinnian Liang, Jian Yang, Hui Huang, Shuangzhi Wu, Peihao Wu, Lu Lu, Zejun Ma, and Zhoujun Li. Enhancing large language model with self-controlled memory framework, 2023. URL https://arxiv.org/abs/2304.13343. 20

Maurice Weber, Daniel Y. Fu, Quentin Anthony, Yonatan Oren, Shane Adams, Anton Alexandrov, Xiaozhong Lyu, Huu Nguyen, Xiaozhe Yao, Virginia Adams, Ben Athiwaratkun, Rahul Chalamala, Kezhen Chen, Max Ryabinin, Tri Dao, Percy Liang, Christopher Ré, Irina Rish, and Ce Zhang. Redpajama: an open dataset for training large language models. *NeurIPS Datasets and Benchmarks Track*, 2024. 22

Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, et al. HuggingFace's Transformers: State-of-the-art natural language processing. *arXiv preprint arXiv:1910.03771*, 2019. 8

Chengxing Xie, Bowen Li, Chang Gao, He Du, Wai Lam, Difan Zou, and Kai Chen. Swe-fixer: Training open-source llms for effective and efficient github issue resolution. *arXiv preprint arXiv:2501.05040*, 2025. 22

Anyi Xu, Bangcai Lin, Bing Xue, Bingxuan Wang, Bingzheng Xu, Bochao Wu, Bowei Zhang, Chaofan Lin, Chen Dong, Chenchen Ling, et al. DeepSeek-V4: Towards highly efficient million-token context intelligence. *arXiv preprint arXiv:2606.19348*, 2026. 8, 21

Peng Xu, Wei Ping, Xianchao Wu, Chejian Xu, Zihan Liu, Mohammad Shoeybi, and Bryan Catanzaro. Chatqa 2: Bridging the gap to proprietary llms in long context and rag capabilities. *arXiv preprint arXiv:2407.14482*, 2024. 22

Wujiang Xu, Zujie Liang, Kai Mei, Hang Gao, Juntao Tan, and Yongfeng Zhang. A-mem: Agentic memory for LLM agents, 2025. URL https://arxiv.org/abs/2502.12110. 20

An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, et al. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*, 2025. 2, 9, 21, 22

Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William Cohen, Ruslan Salakhutdinov, and Christopher D Manning. Hotpotqa: A dataset for diverse, explainable multi-hop question answering. In *Proceedings of the 2018 conference on empirical methods in natural language processing*, pages 2369–2380, 2018. 22

Howard Yen, Tianyu Gao, and Danqi Chen. Long-context language modeling with parallel context encoding. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2588–2610, 2024. 2, 4, 5, 21, 24, 27

Chanwoong Yoon, Taewhoo Lee, Hyeon Hwang, Minbyul Jeong, and Jaewoo Kang. Compact: Compressing retrieved documents actively for question answering. *arXiv preprint arXiv:2407.09014*, 2024. 20

Alex L. Zhang, Tim Kraska, and Omar Khattab. Recursive language models, 2025a. URL https://arxiv.org/abs/2512.24601. 10, 12

Biao Zhang and Rico Sennrich. Root mean square layer normalization. *Advances in neural information processing systems*, 32, 2019. 26

Yanzhao Zhang, Mingxin Li, Dingkun Long, Xin Zhang, Huan Lin, Baosong Yang, Pengjun Xie, An Yang, Dayiheng Liu, Junyang Lin, et al. Qwen3 embedding: Advancing text embedding and reranking through foundation models. *arXiv preprint arXiv:2506.05176*, 2025b. 27

Wenting Zhao, Xiang Ren, Jack Hessel, Claire Cardie, Yejin Choi, and Yuntian Deng. Wildchat: 1m chatgpt interaction logs in the wild. *arXiv preprint arXiv:2405.01470*, 2024. 22

Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Tianle Li, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zhuohan Li, Zi Lin, Eric P Xing, et al. Lmsys-chat-1m: A large-scale real-world llm conversation dataset. *arXiv preprint arXiv:2309.11998*, 2023. 22

Lianmin Zheng, Liangsheng Yin, Zhiqiang Xie, Chuyue Livia Sun, Jeff Huang, Cody Hao Yu, Shiyi Cao, Christos Kozyrakis, Ion Stoica, Joseph E Gonzalez, et al. Sglang: Efficient execution of structured language model programs. *Advances in neural information processing systems*, 37:62557–62583, 2024. 3, 8

Wanjun Zhong, Lianghong Guo, Qiqi Gao, He Ye, and Yanlin Wang. Memorybank: Enhancing large language models with long-term memory, 2023. URL https://arxiv.org/abs/2305.10250. 20

Adam Zweiger, Xinghong Fu, Han Guo, and Yoon Kim. Fast kv compaction via attention matching. *arXiv preprint arXiv:2602.16284*, 2026. 1, 3, 9, 20, 32

Appendix Contents

A Extended Related Work **20**

B Dataset Details **21**
B.1 Continual Pre-training Dataset Curation . . . 21
B.2 Auxiliary Reconstruction Dataset Curation . . . 22
B.3 SFT Dataset Curation . . . 22
B.4 Compression Data Format . . . 23

C Extended Architecture Search **24**
C.1 Prior Soft-Token Methods in Our Design Space . . . 24
C.2 Formal Definitions of Pooling Operators . . . 24
C.3 Adapter Design . . . 26
C.4 Boundary Context . . . 26

D Full-Scale Training Recipe Sweeps **27**
D.1 Pre-trained Encoder Representations . . . 27
D.2 LR Sweep . . . 28
D.3 Decoder Retention Without Compression . . . 28
D.4 Scaling Up Model Size . . . 28
D.5 Packed Sequences . . . 29

E Compute **30**

F Benchmarks and Results **31**
F.1 GSM8K Performance vs Compression Ratio . . . 31
F.2 Baseline Configurations . . . 32
F.3 Timing Implementation . . . 32
F.4 Results . . . 33
F.5 LR Sweep . . . 37
F.6 Architecture Comparison . . . 39
F.7 NIAH Agent Results . . . 43

A Extended Related Work

Hard-token compression. Hard-token methods achieve compression by deleting or rewriting input tokens. Token pruning approaches remove tokens via importance heuristics (Li et al., 2023; Jiang et al., 2023, 2024), while summarization and rewriting methods compress by paraphrasing into fewer tokens (Chuang et al., 2024; Yoon et al., 2024). While widely used in practice (e.g., context compaction in Claude Code and Codex; Anthropic, 2025; OpenAI, 2025) and very effective for reducing prompt length, these methods are intrinsically lossy: once information is dropped or paraphrased in the discrete token space, exact lexical and structural detail is unrecoverable. This limits their ability to support code fidelity, format-sensitive reasoning, and lookup-style in-context learning in long-horizon settings.

Relatedly, many agent systems externalize memory via summarization, retrieval, or structured stores. Examples include hierarchical memory management (Packer et al., 2023), long-term conversational memory (Zhong et al., 2023; Wang et al., 2023), and agentic memory frameworks (Xu et al., 2025). While effective in extending interaction horizons, these approaches typically rely on lossy summaries or task-specific memory representations that are external to the base model.

KV cache compression. KV cache compression reduces the size of the KV cache by evicting entries; most methods rely on a hand-crafted or learned heuristic to select which entries to drop. SnapKV (Li et al., 2024c) compresses the cache by removing entries with low aggregated query attention on a per-head basis. SnapKV can be either prompt-agnostic, where pruning is performed without knowledge of the prompt, or prompt-dependent, where pruning is performed with knowledge of the prompt. Prompt-dependent methods have the downside of requiring labeling of the context and prompt for inputs, and pruning is highly specific to each individual input, limiting multi-turn chat applicability (Li et al., 2024b). To overcome these limitations, self-study (Eyuboglu et al., 2025; Zweiger et al., 2026) can be used to generate synthetic query signals from the context itself, enabling prompt-agnostic pruning without requiring labeled prompt-context pairs.

KVzip (Kim et al., 2025) performs query-agnostic KV cache compression by using teacher-forced context reconstruction as a proxy objective, assigning each KV pair an importance score from the maximum attention it receives during reconstruction, and pruning low-scoring entries. Fast KV Compaction via Attention Matching (Zweiger et al., 2026) compresses the KV cache by fitting a smaller set of keys and values, together with per-entry bias terms, that match the original cache's attention outputs and attention mass on reference queries. Expected Attention (Devoto et al., 2025) predicts KV importance by approximating future queries with a Gaussian distribution, computing each key's expected attention in closed form, and pruning entries with the lowest expected contribution. There are many other KV cache compression techniques and we refer the reader to Awesome KV cache Compression GitHub (Chen, 2023) for a comprehensive review.

We note that KV cache compression methods also pool before eviction: both Li et al. (2024c) and Zweiger et al. (2026) use max pooling, as does Devoto et al. (2025). Li et al. (2024c) further note that the choice of pooling operator does not significantly impact their performance.

It is also possible to do targeted KV cache compression offline for specific prompts. This approach applies to the setting where it is logical to expend more compute to compress the cache once, as this cost will be amortized over many user queries. Synk et al. (2025) compute the full KV cache before applying top-k cache compression and reuse the smaller cache. Eyuboglu et al. (2025) use self-study to learn small KV caches per corpus.

Soft-token compression. A complementary direction encodes long contexts into fewer learned continuous tokens. The most successful examples of soft-token approaches to compression require training on the target context, such as prefix tuning (Li and Liang, 2021; Tan et al., 2024). These methods can achieve comparable in-context performance to the base model with large compression ratios for the specific contexts they are trained on. Other soft-token approaches that do not require training on individual contexts are gist tokens (Mu et al., 2023), in-context autoencoders (Ge et al., 2023) and encoder-decoder compression frameworks (Chevalier et al., 2023; Liao et al., 2025; Dai et al., 2025; Li et al., 2025).

Efficient long sequence modeling. A complementary line of work targets the cost of long-context inference by modifying the underlying sequence model, the attention mechanism, or the positional encoding, rather than compressing the input. Early work scales Transformers to long inputs through segment-level recurrence and external memory: Transformer-XL (Dai et al., 2019) caches hidden states from previous segments to extend the effective receptive field, and the Compressive Transformer (Rae et al., 2019) further compresses older memory entries into a smaller compressed memory.

A second direction replaces softmax attention with formulations that scale linearly in sequence length. Linear-attention Transformers (Katharopoulos et al., 2020) cast attention as a kernel feature map to obtain $O(T)$ inference, state-space models such as S4 (Gu et al., 2021) and Mamba (Gu and Dao, 2024) use recurrent or selective-scan formulations with constant-size hidden state, and more recent hybrids such as Kimi Linear (Kimi Team et al., 2025) interleave linear-attention blocks with sparse full-attention blocks to recover long-range recall while keeping most of the per-layer cost linear. These approaches lower the asymptotic cost of attention but typically trade off exact long-range recall on retrieval-style tasks.

Within the Transformer family, recent DeepSeek models compress the attention computation itself: Multi-head Latent Attention in DeepSeek-V2 (Liu et al., 2024a) projects keys and values into a low-rank latent KV cache, and DeepSeek-V4 (Xu et al., 2026) introduces a hybrid attention architecture that interleaves Compressed Sparse Attention (CSA), which consolidates every m tokens of the KV cache into a single entry and then applies sparse top-k selection over those entries, with Heavily Compressed Attention (HCA), which uses a much larger compression rate $m' \gg m$ but retains dense attention over the resulting entries, jointly reducing both KV memory and per-token attention FLOPs at million-token context lengths.

A separate thread extends the usable context of pre-trained models through their positional encoding: rotary position embeddings (Su et al., 2024) provide relative attention with smooth extrapolation, and YaRN (Peng et al., 2023) rescales RoPE frequencies to extend the context window with minimal additional training.

Finally, parallel context encoding methods such as CEPE (Yen et al., 2024) take a different route: a small encoder produces dense representations of long-context chunks in parallel, and a frozen decoder consumes them through inserted cross-attention layers, exposing long contexts to the decoder without full self-attention over the raw tokens.

These approaches reduce the per-token cost of attention, the cost of positional indexing, or the effective sequence length seen by self-attention, rather than the number of tokens reaching the decoder; they are largely orthogonal to LCLMs and can in principle be composed.

B Dataset Details

B.1 Continual Pre-training Dataset Curation

In early experiments, we only train on reconstruction data and hope the model can learn generalizable patterns. However, we find that training on such data alone allows the model to reconstruct the full sequence with negligible loss, but it generalizes poorly beyond reconstruction-style tasks. In fact, it cannot perform any other tasks, even when the LLM decoder is kept frozen. This might not seem surprising, as we are essentially performing a general version of prefix tuning. Thus, the learned representation, though general for different contexts, collapses to the task of reconstruction. On the other hand, training solely on next-token prediction enables the model to leverage compressed context for downstream generation and reasoning. However, the resulting representations often lack fine-grained fidelity (e.g., exact-string details), leading to weaker reconstruction and brittleness on tasks requiring precise recovery.

Motivated by this trade-off, we train on a mixture of next-token prediction and reconstruction data. Our working hypothesis is that next-token prediction provides task-aligned learning signals that teach the decoder to operate on compressed context, while reconstruction acts as an auxiliary objective that encourages information preservation and can accelerate early-stage representation learning. Theoretically, sufficiently large-scale next-token training alone may recover fine-grained fidelity, but we find that in practice, the reconstruction data serves as an auxiliary data source that accelerates the training of the general compressor.

Curation. Our primary sources are NVIDIA's Nemotron pre-training datasets (NVIDIA et al., 2025; Blakeman et al., 2025). We sample 50M data points from Nemotron's pre-training code and cc code, 50M examples from high-quality and synthetic high-quality cc text, and 50M examples from Nemotron specialized SFT mixtures that cover diverse reasoning traces. A substantial portion of the synthetic data from Nemotron consists of rewrites or generations produced by strong teacher models (e.g., `Qwen3-30B-A3B`, `Qwen3-235B-A22B` (Yang et al., 2025), `GPT-oss-120B` (Agarwal et al., 2025), and `DeepSeek-R1` (Guo et al., 2025)), which alleviate the problem of forgetting. Furthermore, to expose the compressor to longer documents, we sample approximately 500K examples from the OLMo-3 longmino collection (OLMo Team et al., 2025). Finally, to better preserve instruction-following and chat-template behavior, we also include Nemotron pre-training SFT data and convert it into a standardized instruction tuning and multi-turn chat format. Exact dataset composition is reported in Table 1.

Table 1 Dataset mixture for continual pre-training. Pre/Post-Comp are token counts before/after compression. We refer to the unique Nvidia license as "Nvidia"; the license can be found at https://huggingface.co/datasets/nvidia/Nemotron-CC-v2.1/blob/main/LICENSE.md.

Data Name	License	Data Description	Num Samples	Tokens			%
				Pre-Comp	Post-Comp	Trainable	
Nemotron CC	Nvidia	Continual pre-training (CC Text)	84.84M	60.69B	30.38B	30.08B	20.37
Nemotron Code	Nvidia	Continual pre-training (Code)	46.85M	51.18B	25.21B	25.01B	16.90
Nemotron Reasoning	Nvidia	Continual pre-training (Reasoning)	16.62M	52.32B	25.61B	25.50B	17.17
Longmino	odc-by	Continual pre-training (Long-Context)	0.548M	26.62B	4.24B	4.23B	2.84
SFT	cc-by-4.0	SFT with Compressed Prompt	20.67M	42.21B	38.02B	37.54B	25.49
Reconstruction	Nvidia	Reconstruction from Compression	23.53M	50.75B	25.69B	25.18B	17.23
Total	–		**193.06M**	**283.78B**	**149.16B**	**147.54B**	**100.00**

B.2 Auxiliary Reconstruction Dataset Curation

We construct reconstruction data spanning three broad categories: text, code, and LaTeX. Text is sampled from Nemotron CC / math, FineWiki (Penedo, 2025), and RedPajama ArXiv (Weber et al., 2024); code is sampled from RedPajama GitHub, The Stack v2 (we concatenate files from the repo into one document) (Lozhkov et al., 2024), and CodeParrot; and LaTeX is sampled from TexTeller-OCR. Our goal is to maximize data diversity so that the compressor retains sufficient coverage for downstream tasks that may involve any of these modalities.

B.3 SFT Dataset Curation

For reasoning, we curate math, code, and science supervision from the Nemotron SFT datasets and OLMo-3 Dolci-Think (Nathawani et al., 2025a,b; Blakeman et al., 2025; OLMo Team et al., 2025). Unlike the other SFT data, where the prompt or document is compressed, reasoning examples keep the question uncompressed and compress spans of the reference chain of thought; the model is trained to continue reasoning from the compressed spans. About half of the reasoning examples (2.70M) contain compressed chain-of-thought spans, and the rest are left uncompressed.

For long-context instruction tuning, we combine long-context and RAG-style QA datasets (He et al., 2025; Jin et al., 2019; Liu et al., 2024d; Xu et al., 2024; Yang et al., 2018; Liu et al., 2024c; Joshi et al., 2017) with long-context code instruction datasets, including SWE-Fixer (Xie et al., 2025), NextCoder (Aggarwal et al., 2025), PyResBugs (Cotroneo et al., 2025), CommitPack (Muennighoff et al., 2023), and RepoBench (Liu et al., 2023b). Since some public datasets contain outdated or low-quality responses, we regenerate targets by labeling them with `Qwen3-235B-A22B-Instruct-2507`. In addition, we generate three synthetic long-context corpora: (1) repository summarization, using summarization question templates as instructions and concatenated repository files from The Stack v2 (Lozhkov et al., 2024) as the document; (2) document summarization, using long documents from FineWiki (Penedo, 2025) and RedPajama-ArXiv (Weber et al., 2024); and (3) repository code generation, where we hold out one file and ask the model to reconstruct it given the remaining repository context. All synthetic completions are generated by `Qwen3-30B-A3B-Instruct-2507`.

For instruction following, we include multi-turn conversation data from WildChat and LMSYS-Chat-1M (Zhao et al., 2024; Zheng et al., 2023), along with additional alignment-style data obtained by re-labeling with `Qwen3-30B-A3B-Instruct-2507`. We also add OLMo-3 instruction tuning, Nemotron instruction tuning, and the Tulu-3 mixture (Lambert et al., 2024); since these datasets are already generated by strong teacher models, we keep their original responses unchanged. Finally, we retain a small amount of reconstruction data to stabilize training and mitigate drift when optimizing on compressed inputs. For agentic behavior, we include tool-calling, interactive, search, and terminal trajectories from Nemotron-Agentic-v1, Nemotron-SFT-Agentic-v2, and OpenThoughts-Agent-SFT-100K (Blakeman et al., 2025; Raoof et al., 2026), which are kept uncompressed. We also generate synthetic memory-expansion trajectories with Qwen3 models (Yang et al., 2025), in which the source documents are compressed and the model requests individual segments with the `expand` tool before answering from the returned text. Exact dataset composition is reported in Table 2, and token statistics after packing in Table 3.

Table 2 SFT data mixture. Examples are counts after filtering; token statistics per packing pool are in Table 3. The same mixture is used for all compression ratios. Memory-expansion trajectories are generated with Qwen3 models (Apache-2.0); their source documents (e.g., LexGLUE, ACORD, ContractNLI) retain their original licenses. We refer to the unique Nvidia license as "Nvidia"; the license can be found at https://huggingface.co/datasets/nvidia/Nemotron-CC-v2.1/blob/main/LICENSE.md.

Data Name	License	Description	Examples
Reasoning			**5.46M**
Nemotron Reasoning	cc-by-4.0	Nemotron reasoning (compressed CoT)	4.23M
Dolci Think	odc-by	OLMo reasoning (compressed CoT)	1.01M
Nemotron Science	cc-by-4.0	Nemotron science (MCQ, reading QA)	0.22M
Long-Context Instruction Tuning			**7.32M**
Repo Summarization	Apache	Synthetic repo-level summarization	1.35M
Document Summarization	Apache	Synthetic long-document summarization	1.75M
Repo Code Generation	Apache	Synthetic repo-level code generation	2.02M
Long-Context QA	Apache	Long-context / RAG QA	0.84M
Long-Context Code Instruct	Apache	Long-context code instruction tuning	1.36M
General Instruction Following			**5.26M**
Multi-turn Conversation	odc-by	WildChat + LMSYS Chat 1M	2.48M
Dolci Instruct	odc-by	OLMo instruction tuning	1.46M
Tulu3 SFT Mixture	odc-by	Tulu3 instruction tuning	0.84M
Nemotron Instruct	Nvidia	Nemotron instruction tuning	0.35M
HelpSteer	cc-by-4.0	Alignment	0.13M
Reconstruction Data			**2.08M**
Reconstruction	Nvidia	Small reconstruction mixture	2.08M
Agent			**1.35M**
Agent Trajectories	cc-by-4.0 / Apache / MIT	Nemotron-Agentic, OpenThoughts-Agent (uncompressed)	1.19M
Memory Expansion	Apache	Synthetic segment-expansion trajectories (Qwen3)	0.15M
Total			**21.46M**

Table 3 SFT token statistics after packing into sequences of at most 16,384 decoder positions at 16× compression. Decoder positions include latent positions and memory boundary tags; trainable tokens receive next-token prediction loss. Nemotron Science is packed in the other pool.

Pool	Examples	Sequences	Decoder positions	Trainable tokens	Share
Agent	1.35M	0.27M	4.36B	1.10B	11.7%
Reasoning	5.24M	0.71M	11.66B	9.44B	31.4%
Other	14.88M	1.29M	21.16B	17.24B	56.9%
Total	21.46M	2.28M	37.19B	27.78B	100%

B.4 Compression Data Format

Continual pre-training. For the continual pre-training data, given a sequence, we split it into multiple segments, wrap every other segment with the special tokens <|memory_start|> and <|memory_end|>, and treat the wrapped spans as compressed content while leaving the remaining segments uncompressed.

For raw pre-training documents, we construct a fixed alternating pattern of wrapped and unwrapped segments, always beginning with a wrapped segment and ending with an unwrapped segment. Documents longer than 16,384 tokens retain 8,192 trainable (unwrapped) tokens, while shorter documents retain half of their tokens as trainable. The number of segment pairs depends on document length: documents shorter than 2,048 tokens receive 1 pair; documents between 2,048 and 4,096 tokens receive 1-2 pairs; documents between 4,096 and 8,192 tokens receive 2-4 pairs; documents between 8,192 and 16,384 tokens receive 4-6 pairs; and documents longer than 16,384 tokens receive 4-8 pairs. Within each sequence, the first wrapped segment contains 30%-40% of the total wrapped tokens, the last unwrapped segment contains 40%-50% of the total trainable tokens, and the remaining segments are allocated with a $\pm 20\%$ jitter. No segment is allowed to be shorter than 128 tokens.

Instruction tuning. For conversational data, we apply wrapping independently to each message. With 50% probability, the entire message content is enclosed within a single pair of tags. Otherwise, with the remaining 50%

probability, we insert 1-3 wrapped regions inside the message and leave the remaining tokens outside the tags. The number of wrapped regions depends on message length: messages shorter than 32 tokens are always fully wrapped; messages between 32 and 200 tokens receive 1 region; messages between 200 and 600 tokens receive 1-2 regions; and messages of 600 tokens or more receive 1-3 regions. Wrapped regions cover 80%-100% of the message tokens in aggregate, and we enforce a minimum 5% unwrapped gap whenever feasible. If the realized wrapped coverage falls below 70%, we fall back to fully wrapping the message. Since this decision is made independently for each message, multi-turn conversations naturally contain a mixture of fully wrapped turns and turns with small unwrapped spans, with a bias toward high overall coverage.

C Extended Architecture Search

C.1 Prior Soft-Token Methods in Our Design Space

Table 4 describes representative soft-token compressors by the design axes we sweep. Each column corresponds to one of our controlled comparisons: the readout to the pooling operator and encoding granularity (Figure 3, Appendix C.2), the encoder to the attention mask (Figure 3) and initialization (Appendix D.1), the adapter to the adapter comparison (Appendix C.3), and decoder training to the frozen, LoRA, and full-parameter settings (Section 4). These comparisons isolate individual design choices under matched data, scale, and decoder; they do not reproduce the complete prior systems, which also differ in training data, scale, and target tasks.

Table 4 Prior soft-token compressors described by our design axes.

Method	Readout	Encoder	Adapter	Decoder training
AutoCompressor (Chevalier et al., 2023)	learned summary tokens	the decoder LM itself	–	full (OPT) / LoRA (Llama-2)
ICAE (Ge et al., 2023)	learned memory tokens	LLM with LoRA	–	frozen
REFRAG (Lin et al., 2025)	one embedding per chunk ($W = N$)	RoBERTa encoder	projection layer	full, after alignment
E2LLM (Liao et al., 2025)	attention pooling per chunk	pretrained embedding model	adapter	LoRA
GMSA (Tang et al., 2025)	mean over token groups	LLM with LoRA	attention layers	full, in fine-tuning
Simple Context Compression (Feldman and Artzi, 2025)	mean pooling	LLM, bidirectional, LoRA	–	LoRA
xRAG (Cheng et al., 2024)	one retrieval embedding	frozen retriever	2-layer MLP	frozen
CEPE (Yen et al., 2024)	cross-attention to encoder states	small context encoder	cross-attention	frozen
LCLM (ours)	mean pooling, $W = 1024$	embedding model, causal	MLP	full, staged

C.2 Formal Definitions of Pooling Operators

We now give the formal definitions of pooling operators, and illustrate them visually in Figure 8.

Token-based pooling. The most common method is appending one or more learned special pooling tokens per latent output to each encoder window. Let $p_{i,1}, \ldots, p_{i,M_i}$ denote these pooling tokens. We encode

$$(\tilde{h}_1^{(i)}, \ldots, \tilde{h}_{|w_i|+M_i}^{(i)}) = \mathrm{Enc}_\phi([w_i; p_{i,1}; \ldots; p_{i,M_i}]), \tag{2}$$

and use the hidden states at the pooling-token positions:

$$z_k^{(i)} = \tilde{h}_{|w_i|+k}^{(i)}, \qquad k = 1, \ldots, M_i. \tag{3}$$

This generalizes the common EOS-pooling strategy: when $W = N$, a single pooling token is appended for each compression block; when $W > N$, several pooling tokens summarize different regions of the same encoder window.

Mean pooling. For each window w_i, we partition its hidden states into consecutive groups of size N. For latent index $k = 1, \ldots, M_i$, define

$$G_{i,k} = \{(k-1)N + 1, \ldots, \min(kN, |w_i|)\}. \tag{4}$$

The corresponding latent vector is the average hidden state over this group:

$$z_k^{(i)} = \frac{1}{|G_{i,k}|} \sum_{j \in G_{i,k}} h_j^{(i)}. \tag{5}$$

Mean pooling directly aggregates the token-level representations assigned to each compression block. Across our architecture search, mean pooling consistently outperforms EOS-style token pooling, matching observations from prior text-compression and vision-encoder work.

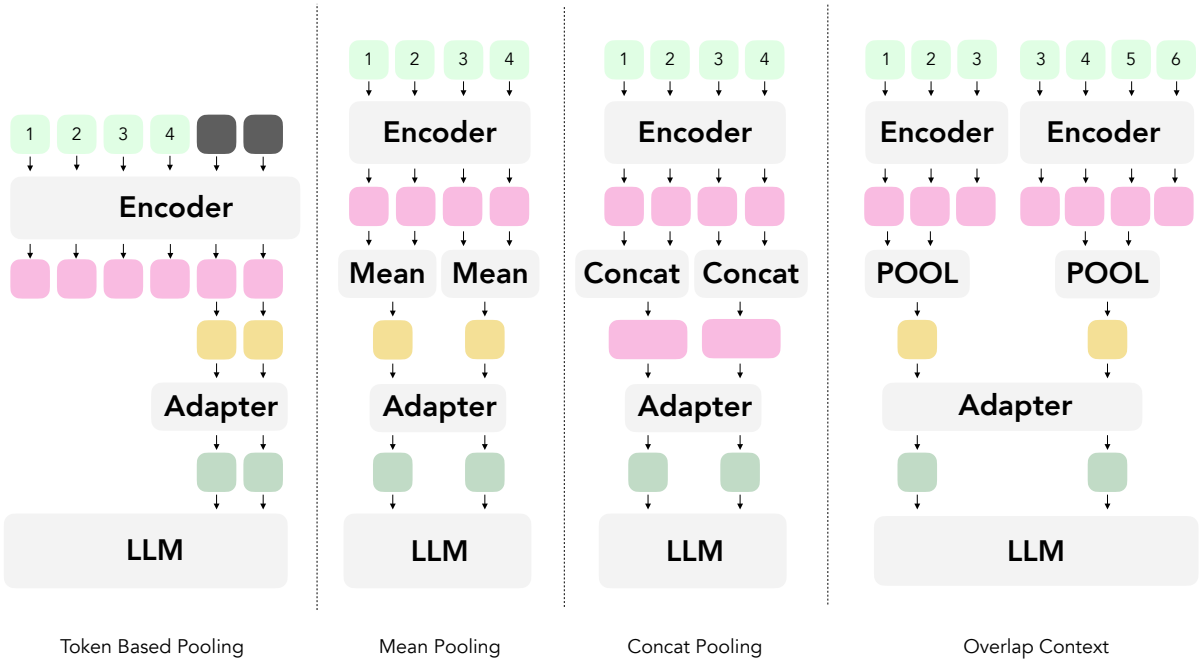


Figure 8 Full architecture design choices for the compressor. From left to right, we visualize token-based pooling, mean pooling, concat pooling, and overlap context.

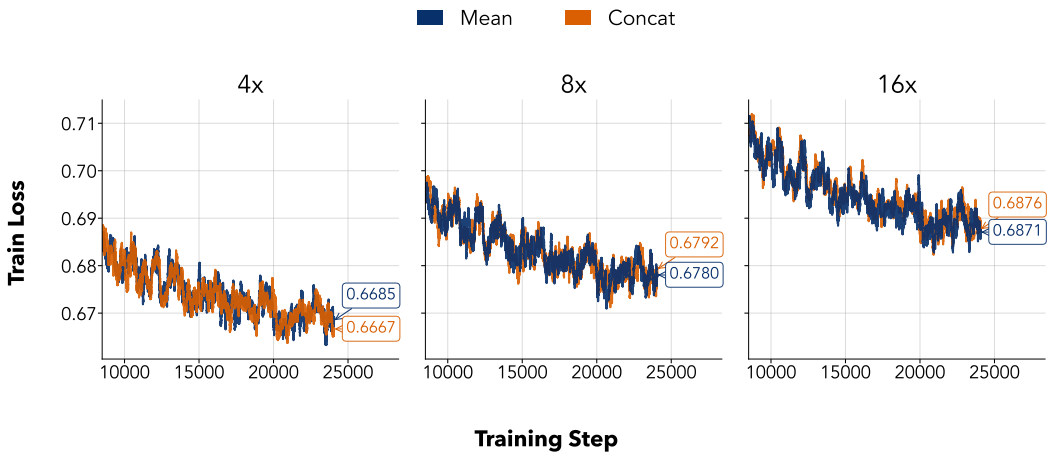


Figure 9 Comparison between mean and concat at scale for all ratios. From Figure 3, we identify mean and concat as the most promising pooling operators, with mean slightly outperforming concat in training loss. We train at larger scale and find that mean remains marginally better than concat at higher compression ratios, but loses to concat at lower compression ratios.

Concat pooling. Rather than averaging or selecting a designated token, concat pooling preserves the full per-token representation by concatenating consecutive encoder hidden states within each compression block. For window w_i and latent index $k = 1, \ldots, M_i$, let $G_{i,k} = \{(k-1)N + 1, \ldots, \min(kN, |w_i|)\}$ as in mean pooling. The concatenated latent is

$$z_k^{(i)} = \left[h_j^{(i)} \right]_{j \in G_{i,k}} \in \mathbb{R}^{N\, d_{\text{enc}}}, \tag{6}$$

i.e., a vector of dimension $N\, d_{\text{enc}}$ formed by stacking the N encoder hidden states. The adapter is widened accordingly so that its first MLP layer takes inputs of dimension $N\, d_{\text{enc}}$ and projects back to the decoder hidden dimension d_{dec}, analogous to the patch-merger / spatial-concat scheme used in Bai et al. (2025).

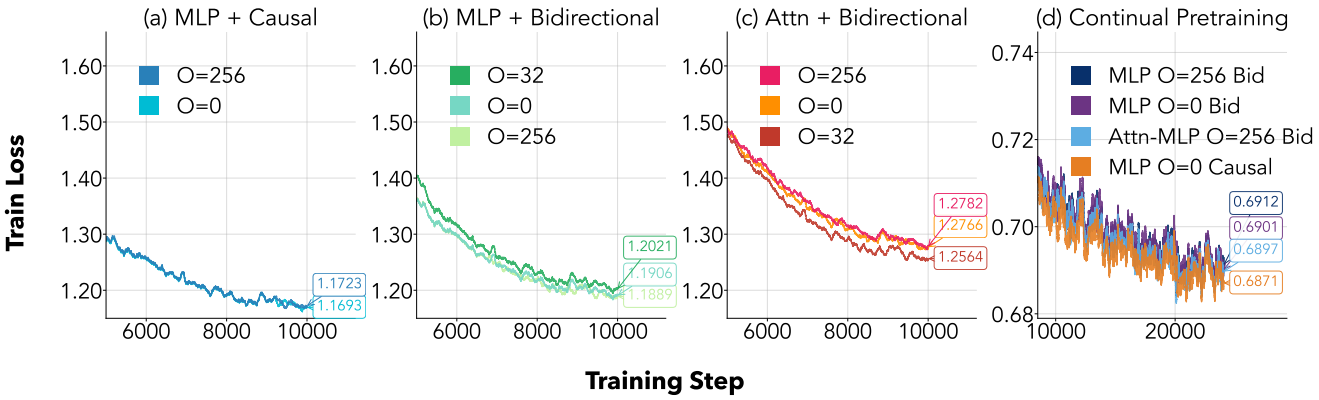


Figure 10 MLP adapters, causal masking, and no boundary overlap give the lowest pre-training loss across the architecture sweep. We sweep combinations of adapter, boundary-overlap, and encoder mask, and find that (1) causal masking is strictly better than bidirectional masking, (2) boundary overlap does not consistently improve over non-overlapping windows while substantially increasing compute, and (3) the MLP adapter outperforms attention-based adapters in the from-scratch runs. At scale (plot d), causal masking and no overlap remain best, while the two adapters are close.

C.3 Adapter Design

The adapter $a(\cdot)$ maps the compressed latent sequence $z_{1:M} \in \mathbb{R}^{M \times d_{\text{enc}}}$ into a sequence $s_{1:M} \in \mathbb{R}^{M \times d_{\text{dec}}}$ in the decoder embedding space. Prior encoder-decoder compression work often assumes $d_{\text{enc}} = d_{\text{dec}}$, but a wide range of encoder-decoder models pair smaller encoders with larger decoders, making the adapter an important design choice. We formally define the two adapter families in our sweep below.

MLP adapter. The MLP adapter is a standard two-layer feed-forward network with a GELU non-linearity (Hendrycks and Gimpel, 2016) and an RMSNorm (Zhang and Sennrich, 2019) pre-normalization, following the projection design used in LLaVA-style VLMs (Liu et al., 2023a):

$$\hat{z}_m = \text{RMSNorm}(z_m), \tag{7}$$

$$s_m = W_2\,\text{GELU}\big(W_1 \hat{z}_m + b_1\big) + b_2, \qquad m = 1, \ldots, M, \tag{8}$$

where $W_1 \in \mathbb{R}^{d_h \times d_{\text{enc}}}$ and $W_2 \in \mathbb{R}^{d_{\text{dec}} \times d_h}$ with hidden dimension d_h. Each latent token is projected independently and the adapter does not mix information across the latent sequence. When the encoder uses concat pooling (Appendix C.2), each z_m has dimension $N\,d_{\text{enc}}$ instead of d_{enc}, so the input dimension of W_1 is widened to $N\,d_{\text{enc}}$ and the first MLP layer is correspondingly $N\times$ larger.

Attention-based adapter. The attention-based adapter adds a single multi-head self-attention layer over the latent sequence before the MLP projection, with an RMSNorm pre-normalization on the latents:

$$\hat{z}_m = \text{RMSNorm}(z_m), \tag{9}$$

$$\tilde{z}_{1:M} = \text{MHSA}\big(\hat{z}_{1:M}\big), \tag{10}$$

$$s_m = W_2\,\text{GELU}\big(W_1 \tilde{z}_m + b_1\big) + b_2, \qquad m = 1, \ldots, M. \tag{11}$$

The self-attention layer operates in the encoder hidden dimension d_{enc} and lets information mix across latent tokens before projection, at the cost of additional parameters and a compute term quadratic in M.

Comparison. Across the from-scratch runs of the adapter sweep in Figure 10, the MLP adapter achieves lower pre-training loss than the attention-based variant while adding less computation, in contrast to prior findings (Tang et al., 2025) that attention-based adapters perform better. We therefore use the MLP adapter as the default in all subsequent experiments.

C.4 Boundary Context

A potential limitation of fixed encoder windows is that information crossing a window boundary is split across two encoder forward passes. Increasing the window size W reduces the frequency of such boundaries, but also

Table 5 Training recipe for compression model for four different stages. We report, for each stage, which module(s) are optimized (adapter / encoder / decoder), peak learning rates, learning-rate schedule, optimizer, global batch size (in tokens), sequence length, training-token budget, and the data source used.

	Stage 0 Adapter Training	Stage 1 Encoder Training	Stage 2 LLM Training	Stage 3 SFT
Training Config				
Adapter Peak LR	1.0×10^{-3}	6.0×10^{-5}	6.0×10^{-5}	3.0×10^{-5}
Encoder Peak LR	NA	6.0×10^{-5}	6.0×10^{-5}	3.0×10^{-5}
LLM Peak LR	NA	NA	1.0×10^{-5}	3.0×10^{-5}
LR scheduler	5% warmup steps with cosine decay to 1.0×10^{-6}			
Optimizer	AdamW ($\beta_1 = 0.9$, $\beta_2 = 0.95$)			
Data				
LLM Batch size (tokens)	4 Million			
LLM Sequence Length	16384			
Encoder Tokens	17.54 B	35.07 B	82.43 B	60.35 B
LLM Tokens (16x)	21.29 B	42.58 B	100.08 B	37.19 B
Total Tokens (16x)	38.83 B	77.65 B	182.51 B	97.54 B
Data Source	pre-training Mixture	pre-training Mixture	pre-training Mixture	SFT

increases compression-time cost. As an alternative, we implement boundary context overlapping, where each encoder window is extended with O neighboring tokens while the number of latent tokens is kept unchanged.

For a target window

$$w_i = x_{(i-1)W+1:\min(iW,T)},$$

we construct an overlapped encoder input $\tilde{w}_i$ by adding neighboring context tokens. For bidirectional attention, we include both left and right overlap when available:

$$\tilde{w}_i = x_{\max(1,(i-1)W-O+1):\min(iW+O,T)}. \tag{12}$$

We provide a visualization in the rightmost figure in Figure 8. For causal attention, future context is unavailable, so we include only left overlap:

$$\tilde{w}_i = x_{\max(1,(i-1)W-O+1):\min(iW,T)}. \tag{13}$$

After encoding $\tilde{w}_i$, we apply pooling only to hidden states corresponding to the original non-overlapped region w_i, and discard pooled outputs that correspond to overlap tokens. Thus, overlap changes the encoder context available to each window but does not change the number of latent tokens passed to the decoder.

We sweep $O \in \{0, 32, 256\}$ for $W \in \{256, 1024\}$ in Figure 10. Overlap does not noticeably improve pre-training loss in most settings, while increasing training cost because boundary tokens are encoded multiple times. For example, with $W = 1024$ and $O = 256$, each interior bidirectional window processes 1536 tokens instead of 1024. We therefore do not use overlap in the default architecture.

D Full-Scale Training Recipe Sweeps

D.1 Pre-trained Encoder Representations

For the encoder initialization, to determine whether an embedding model's representation (Lin et al., 2025; Liao et al., 2025; Yen et al., 2024) is better than that of a standard language model (Tang et al., 2025; Chevalier et al., 2023; Feldman and Artzi, 2025), we conduct an experiment with the EOS pooling token. We can see from Figure 11 (d) that using the embedding model as the encoder outperforms using an LLM as the encoder. Note the embedding model was initialized from the LLM (Zhang et al., 2025b). This shows that embedding model training provides a better representation for the encoder, though the gap gradually narrows with longer training. We also show downstream evaluation performance in Table 35.

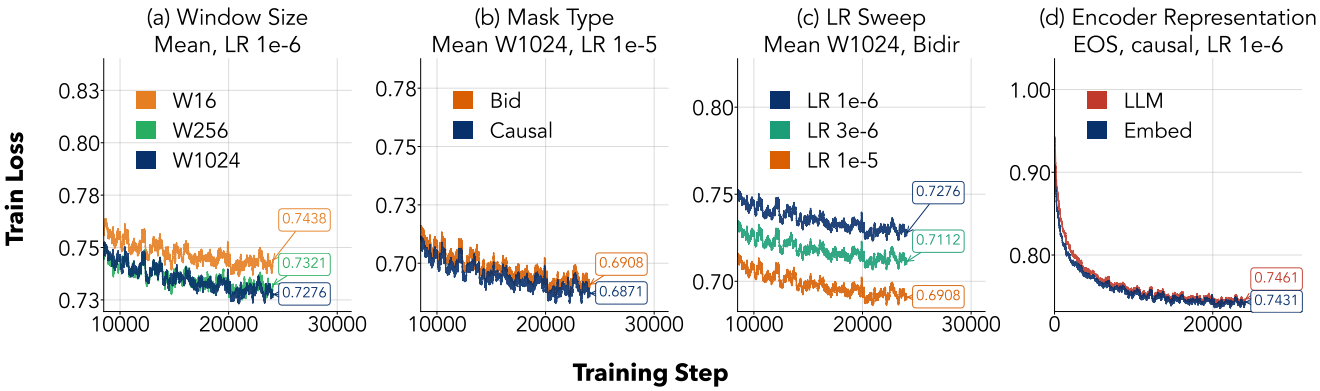


Figure 11 Architecture choices selected by the small-scale sweep continue to hold at scale. Continual pre-training loss for the full-pipeline runs at compression ratio 16×. Across (a) encoder window size, (b) encoder attention mask, (c) stage-2 learning rate, and (d) encoder initialization (LM vs. embedding), the configuration chosen from the from-scratch sweep ($W = 1024$, causal masking, embedding-initialized encoder) yields the lowest pre-training loss.

D.2 LR Sweep

Language models are prone to catastrophic forgetting during continual pre-training and SFT (Luo et al., 2025; Li et al., 2024a). To actively avoid forgetting, we sweep the peak learning rate for stages 2 and 3 of training. To test for catastrophic forgetting, we then benchmark the language model without compression after SFT and check that its performance is comparable to the benchmark performance of the language model before training. For stage 2, we sweep from 1×10^{-6} to 1×10^{-5} and find that a larger learning rate improves downstream benchmarks with minimal drift from the original model (Appendix D.3). For stage 3, we sweep from 2×10^{-5} to 5×10^{-5}. We do not observe significant differences in those SFT LRs across different models; we report the detailed numbers on downstream benchmarks in Appendix F.5.

D.3 Decoder Retention Without Compression

To measure how much training changes the decoder itself, we evaluate our 4×, 8×, and 16× models with compression disabled, so that the decoder reads the full context as ordinary tokens, and compare them with the original `Qwen3-4B-Instruct-2507`. All four models use identical prompts, output budgets, and scoring; the base model is re-evaluated under this protocol. As shown in Table 6, the trained decoders stay within 1.7 points of the original model on RULER at 4k, within 2.8 points at 8k, and within 4.5 points at 16k, and slightly exceed it on LongBench.

Table 6 Accuracy with compression disabled. The decoders of our 4×, 8×, and 16× models read the uncompressed context and are compared with the original model under the same protocol (RULER: 13 tasks, 6,500 questions per context length; LongBench EN16: macro average over 16 English and code tasks). Differences from the original model are in parentheses.

	`Qwen3-4B-Instruct-2507`	LCLM 4× decoder	LCLM 8× decoder	LCLM 16× decoder
RULER 4k	94.49	94.27 (−0.22)	92.86 (−1.63)	93.41 (−1.08)
RULER 8k	93.44	91.55 (−1.89)	90.73 (−2.71)	90.98 (−2.46)
RULER 16k	93.53	91.54 (−1.99)	89.46 (−4.07)	89.10 (−4.43)
LongBench EN16	45.30	47.13 (+1.83)	45.98 (+0.68)	46.91 (+1.61)

D.4 Scaling Up Model Size

To study scaling behavior, we ask two questions: (i) does a larger encoder improve compression quality, and (ii) does a larger decoder, which has a wider hidden dimension that could compress more information per token, improve performance under compression? Accordingly, we add two model configurations: (1) using the larger `Qwen3-Embedding-4B` as the encoder with the same size `Qwen3-4B-Instruct-2507` decoder; and (2) using the same size `Qwen3-Embedding-0.6B` encoder and the larger `Qwen3-8B` as the decoder. In Figure 12 (a), we see that increasing the size of the decoder is much more beneficial in terms of pre-training loss than increasing the size of the encoder.

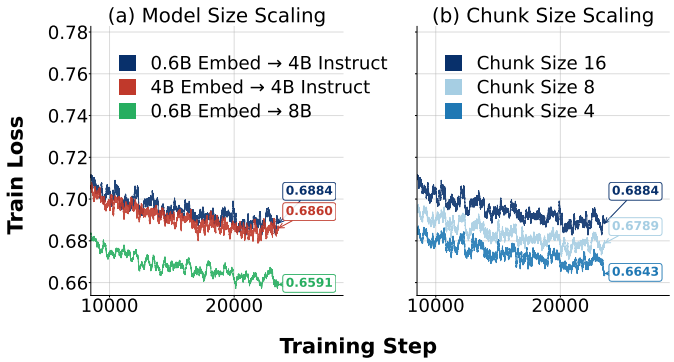


Figure 12 **Scaling the decoder helps more than scaling the encoder, and lower compression ratios yield lower pre-training loss. Left:** pre-training loss at the end of stage 2 with compression ratio 16× across encoder/decoder size combinations; increasing the decoder size produces a much larger drop in loss than increasing the encoder size. **Right:** pre-training loss at the end of stage 2 for our 0.6B-4B models across compression ratios; lower compression ratios yield lower pre-training loss. Note that the 4B-encoder and 8B-decoder runs use 256 ranks rather than 128 ranks. Because our stateful Parquet dataloader assigns work by rank, these runs complete slightly fewer optimization steps. For a clean comparison, we therefore truncate all loss curves at the same optimization step in the plot above.

However, the scaling results in Table 7 are mixed. The 0.6B encoder performs best across the RULER tasks, while the 4B encoder performs best on the remaining evaluations. In contrast, the 8B decoder achieves substantially lower pre-training loss but does not yield the downstream gains we expected. We suspect this is partly due to a mismatch between the training mixture and the decoder initialization: our data curation and recipe were tuned around the 4B instruct decoder, whereas the 8B decoder is a hybrid-thinking model and may require a different data distribution. Moreover, the 4B instruct model already outperforms the 8B model in the full-cache setting on several evaluations, which may limit the apparent benefit of scaling the decoder in our current setup.

Table 7 Encoder/decoder scaling at 16× compression. These models were trained with an earlier version of our SFT data, so the 0.6B/4B row differs from Table 9; LongBench averages all 21 tasks.

Enc / Dec	RULER 4K	RULER 8K	RULER 16K	LongBench	LongHealth5	GSM8K
0.6B / 4B	**75.06**	**70.23**	**65.91**	37.33	67.50	81.05
0.6B / 8B	71.20	67.40	62.22	37.80	64.80	77.30
4B / 4B	71.30	66.90	62.00	**39.00**	**69.80**	**83.20**

D.5 Packed Sequences

Since our corpus contains many long-context examples with highly variable lengths, conventional batching would waste substantial computation on padding. To improve training throughput, we pack multiple examples into a single sequence. Because we continually pre-train an instruction-tuned model, we prevent cross-example information leakage by resetting the attention mask at example boundaries, yielding a block-diagonal attention pattern. We implement this efficiently using variable-length attention kernels (Dao, 2023; Dong et al., 2024).

E Compute

All pre-training-from-scratch experiments are run on 16 nodes of MI300A GPUs, each containing 4 GPUs. Full-scale continual pre-training experiments for 0.6B-4B models are trained on 32 nodes of MI300A GPUs, each containing 4 GPUs.

All evaluations are conducted on H200 GPUs.

All scaling experiments (`Qwen3-Embedding-4B` encoder, `Qwen3-8B` decoder) are trained on H200 clusters with 32 nodes to alleviate memory issues.

F Benchmarks and Results

Table 8 Evaluation benchmarks and what content is compressed. We evaluate on long-context suites and on GSM8K, a short and information-dense benchmark. The "What is Compressed" column specifies which part of the input is placed inside the compressible memory block, while remaining components, such as instructions, are kept as hard tokens when applicable.

Benchmark	Description	What is Compressed
Long-context		
RULER	Synthetic long-context stress tests from 4K–16K.	Task-specific context
LongBench	Aggregated long-context suite covering single- and multi-document QA, summarization, few-shot learning, synthetic tasks, and code completion.	Task-specific context
LongHealth	Clinical-document multiple-choice QA with patient-case context. Five documents are concatenated. No CoT prompting.	Task-specific context
Fine-Grained Compression		
GSM8K	Short and information-dense grade-school math word problems.	Whole prompt

F.1 GSM8K Performance vs Compression Ratio

The long-context benchmarks show that LCLMs can compress extended contexts efficiently. We next evaluate on GSM8K (Cobbe et al., 2021) to test whether LCLMs can also handle short, information-dense inputs, where nearly every token may be relevant to the answer. In Figure 13, LCLMs achieve the highest accuracy across all compression ratios on GSM8K, with particularly strong gains over baselines at higher compression ratios. These are aggressive compression settings: compression ratios of 16× and 8× remove 93.75% and 87.5% of the input tokens, respectively. Strong performance in this setting demonstrates that LCLMs are not specialized only for long-context QA; they can also compress short, dense contexts, making them practical for general-purpose deployment.

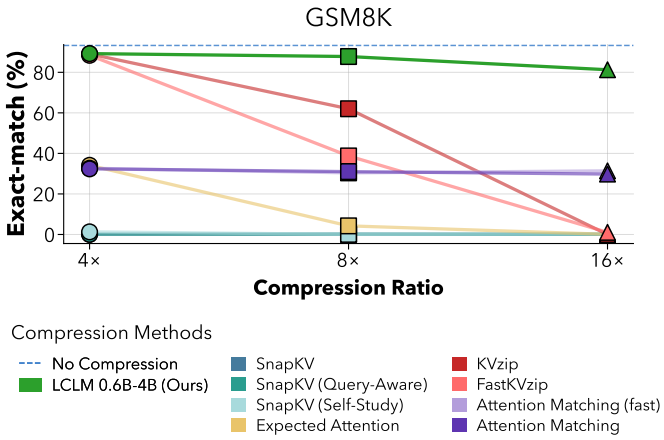


Figure 13 Latent Context Language Models can compress small dense contexts with high accuracy. We plot GSM8K accuracy over different compression ratios and find that LCLMs maintain much higher accuracy at larger compression ratios. Full data are provided in Table 9. We do not report compression time for GSM8K because the context is small enough that compression-time measurements are too noisy to support reliable conclusions.

F.2 Baseline Configurations

Query-aware methods, such as SnapKV, are designed to compress a context given a known query; however, the resulting compressed cache may not generalize well to unseen queries or future turns. Query-agnostic methods, such as KVzip and FastKVzip, compress a cache without access to the downstream query, often using reconstruction-style objectives such as repeating the original context. Self-study, as introduced by Eyuboglu et al. (2025), performs compression using synthetic queries to expand coverage over potential inference-time queries. These synthetic queries are typically produced by prompting the target decoder to generate questions about the context. In Attention Matching (Zweiger et al., 2026), the authors experiment with a mixed query-agnostic and self-study configuration and find that self-study has limited impact on the repeat-prefill prompt. Since self-study introduces additional compression-time overhead, we do not apply it to RULER.

It is also important to distinguish algorithmic and theoretical cache reduction from systems-realizable cache reduction. Some KV cache methods produce an actual smaller cache that can, in principle, reduce decoding memory and latency. Others, such as Kim et al. (2025, 2026), are evaluated by masking or subsampling entries from a full cache, or by using non-uniform eviction patterns across heads and layers. These approaches are useful for measuring the quality of a compressed cache, but they do not necessarily translate into wall-clock speedups or memory savings in standard inference engines without specialized kernels that have not been released or even implemented.

F.3 Timing Implementation

Comparison with KV cache compression. In Figures 1 and 4, the decoder of every method, including ours, runs in Hugging Face Transformers with fused attention kernels: FlashAttention-2 (Dao, 2023) for the baselines and PyTorch scaled dot-product attention for ours. Our encoder additionally uses FlashAttention, `torch.compile`, and CUDA graphs, with an encoder batch size of 32 windows.

Comparison with uncompressed models. In Figure 5, the decoder of every model, including ours, is served with vLLM (Kwon et al., 2023) in BF16 on a single H200, with batch size 1, CUDA graphs enabled, and prefix caching disabled. Our encoder uses the same optimized implementation as above. For LCLMs, we time the encoder and the vLLM decoder prefill over a sequence of the same length as the compressed latent sequence, and report their sum. Each point is the fastest of 20 timed runs after 30 warmup runs.

F.4 Results

Table 9 Summary across compression ratios: RULER, LongBench, LongHealth5, GSM8K. Rows are grouped by compression ratio (16x, 8x, 4x). Best per column within each ratio block in **bold**, second-best underlined.

Model	RULER 4k	RULER 8k	RULER 16k	LB en16	LB cn5	LongHealth5	GSM8K
Qwen 4B Instruct (full KV)	94.49	93.44	93.53	45.30	46.54	76.75	93.25
16x compression							
LCLM 0.6b-4b (Mean)	**77.01**	**72.10**	**67.04**	38.26	31.38	66.25	**81.27**
ExpAttn	40.97	46.93	50.67	29.44	20.51	30.25	0.08
SnapKV	14.31	14.11	18.27	26.96	18.30	10.50	0.08
SnapKV-QA	20.54	39.09	61.96	38.52	**42.26**	**76.25**	0.08
KVzip	62.73	61.73	57.97	25.86	17.12	67.00	0.00
KVzipFast	40.01	44.70	47.86	25.76	17.08	18.25	1.06
AM-Fast	53.09	55.45	37.17	**40.49**	31.51	70.50	31.39
AM-Slow	69.21	70.38	42.02	40.26	31.85	72.75	29.80
8x compression							
LCLM 0.6b-4b (Mean)	89.19	88.06	85.76	41.74	36.65	70.75	**87.79**
ExpAttn	65.08	65.28	65.41	38.13	30.56	51.50	4.25
SnapKV	19.64	19.42	25.39	32.28	25.18	13.75	0.15
SnapKV-QA	47.66	60.32	72.10	41.73	**45.41**	**76.00**	0.15
KVzip	**90.27**	**88.34**	**88.65**	40.72	36.09	72.75	62.02
KVzipFast	86.08	84.59	85.27	38.17	30.91	68.50	38.74
AM-Fast	89.63	87.05	56.06	**44.50**	41.52	74.50	30.40
AM-Slow	89.20	87.05	59.34	44.22	41.11	74.75	31.01
4x compression							
LCLM 0.6b-4b (Mean)	93.44	92.49	90.93	45.57	39.71	74.75	**89.23**
ExpAttn	85.26	83.09	82.51	41.95	41.57	66.00	34.12
SnapKV	27.84	31.02	36.33	38.18	34.22	26.00	0.08
SnapKV-QA	69.24	76.16	81.06	43.74	46.19	75.75	0.08
KVzip	**94.43**	**93.28**	**93.93**	45.30	47.45	76.00	89.08
KVzipFast	94.17	93.15	93.64	45.16	**48.66**	**76.25**	88.40
AM-Fast	92.69	91.60	82.37	45.48	48.39	74.50	32.68
AM-Slow	92.52	91.47	81.98	**45.78**	48.07	73.50	32.37

Table 10 RULER per-task at 4k context. Rows grouped by compression ratio (16x, 8x, 4x).

Model	ns1	ns2	ns3	nm1	nm2	nm3	nmv	nmq	vt	cwe	fwe	qa1	qa2	AVG
Qwen 4B Instruct	100.00	100.00	99.80	99.80	96.80	100.00	99.60	100.00	99.60	97.72	87.60	84.00	63.40	94.49
16x compression														
LCLM 0.6b-4b (Mean)	96.60	90.20	58.40	**86.00**	84.20	**63.20**	**74.85**	**75.55**	82.92	76.04	**92.80**	67.60	**52.80**	**77.01**
ExpAttn	**100.00**	49.20	0.60	57.60	2.20	0.00	56.75	25.70	97.80	18.36	66.20	33.20	25.00	40.97
SnapKV	4.00	1.60	2.40	9.80	1.80	0.00	8.55	8.55	5.48	25.30	73.13	24.80	20.60	14.31
SnapKV-QA	16.60	7.20	2.40	10.20	4.80	0.00	9.10	10.50	15.80	25.14	57.67	60.60	47.00	20.54
KVzip	**100.00**	91.60	**96.00**	62.40	69.80	41.00	32.85	43.35	**99.96**	21.80	70.13	49.00	37.60	62.73
KVzipFast	**100.00**	38.80	31.60	23.00	65.40	26.00	7.70	9.80	99.60	12.66	54.73	24.80	26.00	40.01
AM-Fast	**100.00**	52.80	7.80	33.80	89.00	3.80	22.60	17.85	95.16	**76.82**	81.40	63.20	46.00	53.09
AM-Slow	**100.00**	**95.20**	31.60	74.00	**94.00**	9.40	65.40	56.50	94.96	74.52	81.40	**70.60**	52.20	69.21
8x compression														
LCLM 0.6b-4b (Mean)	98.60	98.40	77.80	95.60	95.40	89.80	96.40	94.90	94.72	**92.28**	**93.20**	75.20	57.20	89.19
ExpAttn	**100.00**	96.20	5.20	92.20	69.20	0.40	93.00	80.20	**100.00**	48.08	85.20	39.20	37.20	65.08
SnapKV	12.40	8.20	2.40	12.60	3.20	0.00	11.20	12.60	9.80	49.06	80.07	31.00	22.80	19.64
SnapKV-QA	54.40	97.40	2.40	46.80	14.40	2.20	21.25	79.65	52.60	43.22	70.20	78.20	56.80	47.66
KVzip	**100.00**	98.20	**99.80**	97.20	98.60	99.40	82.85	96.70	99.96	77.48	82.67	**79.60**	**61.00**	**90.27**
KVzipFast	**100.00**	**100.00**	**99.80**	91.80	98.20	**99.60**	54.15	93.70	99.96	66.52	81.27	76.40	57.60	86.08
AM-Fast	**100.00**	99.40	93.20	**99.20**	98.20	93.00	96.90	**98.45**	96.60	65.90	84.53	79.40	60.40	89.63
AM-Slow	99.80	99.40	90.60	98.20	**98.80**	93.60	**97.65**	97.35	97.04	64.80	84.13	78.60	59.60	89.20
4x compression														
LCLM 0.6b-4b (Mean)	99.60	99.40	97.80	99.40	**99.80**	98.20	**99.40**	99.50	96.76	90.76	**90.93**	79.40	**63.80**	93.44
ExpAttn	**100.00**	99.80	66.60	**100.00**	99.40	53.60	98.15	99.65	**100.00**	86.74	87.07	62.80	54.60	85.26
SnapKV	32.00	14.80	2.60	18.60	6.00	0.00	13.80	16.55	27.00	71.72	81.87	47.60	29.40	27.84
SnapKV-QA	95.00	**100.00**	2.60	98.60	30.20	40.40	54.30	99.35	85.04	72.14	78.73	81.60	62.20	69.24
KVzip	**100.00**	99.40	**100.00**	**100.00**	99.60	99.80	99.25	99.80	99.96	**96.24**	85.80	**85.40**	62.40	**94.43**
KVzipFast	**100.00**	**100.00**	**100.00**	**100.00**	99.60	99.80	98.50	**100.00**	99.96	95.96	84.40	83.00	63.00	94.17
AM-Fast	**100.00**	99.80	96.80	98.80	99.60	**100.00**	98.75	99.40	99.08	78.28	89.40	81.80	63.20	92.69
AM-Slow	**100.00**	**100.00**	97.60	99.00	99.60	99.60	99.10	99.20	98.56	76.98	88.07	81.80	63.20	92.52

Table 11 RULER per-task at 8k context. Rows grouped by compression ratio (16x, 8x, 4x).

Model	ns1	ns2	ns3	nm1	nm2	nm3	nmv	nmq	vt	cwe	fwe	qa1	qa2	AVG
Qwen 4B Instruct	100.00	100.00	99.80	99.40	96.40	99.20	98.05	100.00	99.52	94.98	87.33	81.20	58.80	93.44
16x compression														
LCLM 0.6b-4b (Mean)	96.00	90.40	61.80	**84.20**	74.20	44.20	71.55	**77.70**	77.28	52.38	**93.33**	64.40	49.80	**72.10**
ExpAttn	**100.00**	76.60	1.00	81.20	2.40	0.00	69.30	55.85	99.08	10.52	60.73	30.40	23.00	46.93
SnapKV	8.60	1.20	2.40	10.00	0.80	0.00	9.80	9.80	5.04	24.78	78.60	15.40	17.00	14.11
SnapKV-QA	88.20	62.00	2.40	14.60	9.00	7.60	10.30	44.25	52.68	21.16	68.13	**74.20**	**53.60**	39.09
KVzip	**100.00**	**98.00**	**95.80**	64.80	56.40	**45.20**	31.25	51.05	**100.00**	8.16	73.80	43.20	34.80	61.73
KVzipFast	**100.00**	43.60	38.80	30.80	65.20	39.00	10.85	17.15	99.96	7.42	77.73	23.80	26.80	44.70
AM-Fast	**100.00**	60.00	11.80	42.00	91.24	8.80	28.55	31.70	93.56	**72.90**	73.47	62.80	44.00	55.45
AM-Slow	99.20	94.80	38.40	79.40	**93.20**	11.60	**73.45**	72.95	94.08	71.22	75.07	65.20	46.40	70.38
8x compression														
LCLM 0.6b-4b (Mean)	99.00	98.00	85.20	95.20	94.00	89.40	95.45	94.45	93.20	**78.84**	**92.60**	74.40	55.00	88.06
ExpAttn	**100.00**	98.20	6.20	97.00	69.60	0.20	91.95	94.70	**100.00**	34.12	78.73	42.60	35.40	65.28
SnapKV	37.80	2.40	2.40	11.60	1.60	0.00	10.40	10.50	19.56	34.82	86.33	16.80	18.20	19.42
SnapKV-QA	99.60	96.60	2.40	63.20	22.00	40.40	24.85	97.60	91.44	35.48	75.33	**77.40**	57.80	60.32
KVzip	**100.00**	99.40	**99.80**	88.80	97.60	98.00	82.25	**98.60**	**100.00**	72.24	83.53	71.40	56.80	**88.34**
KVzipFast	**100.00**	**99.80**	**99.80**	79.80	**98.60**	**99.00**	59.50	95.50	**100.00**	58.46	83.67	67.40	**58.20**	84.59
AM-Fast	97.60	98.60	91.40	**98.60**	98.20	95.00	**96.05**	97.20	93.84	54.00	79.60	74.20	57.40	87.05
AM-Slow	96.80	99.40	93.40	96.40	98.40	95.20	95.70	95.84	94.28	52.88	80.60	75.00	57.80	87.05
4x compression														
LCLM 0.6b-4b (Mean)	99.60	99.60	96.80	99.00	99.20	96.60	99.10	99.20	96.00	81.68	**94.33**	**79.40**	**61.80**	92.49
ExpAttn	**100.00**	**100.00**	60.60	**99.20**	98.60	48.60	97.65	99.80	**100.00**	77.74	85.13	60.80	52.00	83.09
SnapKV	67.00	23.60	2.40	18.80	3.80	1.00	18.20	17.90	47.52	60.22	92.87	23.00	27.00	31.02
SnapKV-QA	**100.00**	99.80	2.40	99.00	46.80	87.40	66.75	99.95	99.28	66.00	86.27	77.00	59.40	76.16
KVzip	**100.00**	**100.00**	**100.00**	98.40	**100.00**	99.80	**99.20**	**100.00**	**100.00**	**91.86**	88.40	77.20	57.80	**93.28**
KVzipFast	**100.00**	**100.00**	**100.00**	98.40	**100.00**	**100.00**	98.75	99.95	**100.00**	91.84	86.27	75.60	60.20	93.15
AM-Fast	99.40	**100.00**	98.00	98.60	99.40	99.20	96.40	98.90	96.96	78.62	86.47	79.20	59.60	91.60
AM-Slow	99.60	**100.00**	98.40	98.40	99.40	98.20	95.95	98.05	97.40	78.02	87.33	78.60	59.80	91.47

Table 12 RULER per-task at 16k context. Rows grouped by compression ratio (16x, 8x, 4x).

Model	ns1	ns2	ns3	nm1	nm2	nm3	nmv	nmq	vt	cwe	fwe	qa1	qa2	AVG
Qwen 4B Instruct	100.00	100.00	99.80	99.20	94.80	99.80	98.55	99.95	99.84	87.90	98.80	78.20	59.00	93.53
16x compression														
LCLM 0.6b-4b (Mean)	97.00	90.00	59.40	78.20	**65.20**	34.40	71.80	70.30	72.28	21.72	**99.20**	62.60	49.40	**67.04**
ExpAttn	**100.00**	83.00	1.60	86.60	1.60	0.00	**78.30**	76.35	99.32	8.52	71.07	28.40	24.00	50.67
SnapKV	41.80	4.20	2.40	11.20	2.00	0.20	10.00	10.85	17.60	16.64	92.20	12.40	16.00	18.27
SnapKV-QA	99.80	**99.40**	2.40	**88.40**	11.60	27.00	30.75	**99.15**	93.68	22.22	98.33	**74.80**	**58.00**	61.96
KVzip	**100.00**	92.20	**91.60**	55.20	37.60	18.40	35.70	53.10	**99.96**	3.18	88.87	41.20	36.60	57.97
KVzipFast	**100.00**	53.60	42.00	38.40	64.40	**44.00**	19.05	25.40	99.76	3.58	79.00	23.40	29.60	47.86
AM-Fast	52.05	33.00	0.20	28.40	20.60	0.00	19.40	23.70	54.92	**55.44**	94.67	58.60	42.20	37.17
AM-Slow	61.00	31.20	1.20	38.60	24.80	0.00	32.30	34.50	66.84	52.28	93.33	63.00	47.20	42.02
8x compression														
LCLM 0.6b-4b (Mean)	98.60	98.20	83.60	94.80	87.40	83.60	93.25	92.85	93.68	62.32	98.80	74.00	53.80	85.76
ExpAttn	**100.00**	97.80	6.80	97.00	58.00	0.60	**93.65**	96.90	99.92	30.14	95.73	39.00	34.80	65.41
SnapKV	63.60	11.00	2.40	12.40	2.60	0.60	11.70	12.65	47.72	34.86	97.40	14.40	18.80	25.39
SnapKV-QA	**100.00**	**100.00**	2.40	**99.20**	27.40	67.60	65.70	**99.90**	99.00	40.72	**99.53**	**75.40**	**60.40**	72.10
KVzip	**100.00**	99.80	99.60	90.40	92.60	93.00	83.95	98.65	**99.96**	66.62	98.67	71.40	57.80	**88.65**
KVzipFast	**100.00**	**100.00**	**100.00**	76.80	**95.60**	**96.00**	65.85	96.95	**99.96**	58.10	97.07	68.40	53.80	85.27
AM-Fast	51.20	60.20	10.40	58.20	75.40	13.40	56.45	58.20	55.24	**67.10**	96.20	74.00	52.80	56.06
AM-Slow	53.80	62.20	13.00	70.60	78.60	17.80	63.25	66.80	60.84	65.30	95.20	71.80	52.20	59.34
4x compression														
LCLM 0.6b-4b (Mean)	99.40	99.80	98.00	99.00	99.60	95.20	98.65	98.30	94.44	64.06	99.07	76.80	59.80	90.93
ExpAttn	**100.00**	99.80	57.40	99.00	98.00	34.60	98.95	99.60	**99.96**	73.42	98.47	61.20	52.20	82.51
SnapKV	80.60	46.40	2.40	19.20	4.00	2.20	14.95	18.30	78.16	59.58	99.13	22.00	25.40	36.33
SnapKV-QA	**100.00**	99.80	6.20	**99.80**	59.00	94.60	93.45	**99.95**	99.88	66.60	**99.73**	75.40	59.40	81.06
KVzip	**100.00**	**100.00**	**100.00**	99.20	**99.80**	**100.00**	**99.60**	**99.95**	**99.96**	85.00	98.80	77.60	**61.20**	**93.93**
KVzipFast	**100.00**	**100.00**	**100.00**	98.60	99.40	99.80	99.20	**99.95**	**99.96**	**85.46**	98.93	76.00	60.00	93.64
AM-Fast	60.00	94.40	80.20	93.00	91.60	91.20	91.60	92.45	63.56	79.80	96.46	**78.80**	57.80	82.37
AM-Slow	54.40	94.60	78.60	95.00	94.00	92.00	91.85	94.10	63.16	79.08	95.40	77.60	56.00	81.98

Table 13 LongBench English (en16) per-subtask. Rows grouped by compression ratio (16x, 8x, 4x).

Model	narr	qasp	mfq	hpot	2wiki	musq	govr	qmsm	mnews	trec	triv	samsm	pcnt	pr	lcc	rbp	AVG
Qwen 4B Instruct	30.68	44.05	47.39	50.55	36.25	16.39	30.56	22.53	23.63	76.50	85.09	40.06	3.44	94.98	62.43	60.32	45.30
16x compression																	
LCLM 0.6b-4b (Mean)	26.12	40.43	42.38	46.55	**37.30**	24.24	28.96	23.18	23.45	54.00	**86.58**	33.80	2.66	67.25	37.76	37.54	38.26
ExpAttn	17.67	23.16	28.88	33.76	23.45	8.78	23.06	19.86	19.54	36.00	81.02	31.80	7.36	30.32	28.24	**58.14**	29.44
SnapKV	12.97	13.68	22.14	20.93	16.13	5.61	21.61	17.66	14.70	32.75	85.87	36.57	3.25	17.22	**53.42**	56.78	26.96
SnapKV-QA	**28.36**	28.81	40.37	**53.20**	29.33	17.33	21.57	22.78	14.06	35.00	83.92	37.66	2.37	**92.50**	**53.42**	55.60	38.52
KVzip	16.08	16.38	32.53	10.30	8.68	4.93	24.43	21.38	17.49	43.00	78.22	31.34	6.50	14.00	31.96	56.53	25.86
KVzipFast	11.96	18.26	24.60	23.57	18.08	4.31	20.15	20.48	15.03	34.00	72.78	34.22	**7.44**	8.25	46.89	52.12	25.76
AM-Fast	26.31	**42.64**	45.22	48.31	30.57	24.17	31.59	23.90	**24.14**	**74.50**	85.84	**40.37**	5.11	47.04	48.26	49.81	**40.49**
AM-Slow	24.83	41.04	**45.28**	45.85	30.76	**24.84**	**31.62**	**24.02**	24.08	74.00	86.17	39.60	5.08	47.38	48.79	50.76	40.26
8x compression																	
LCLM 0.6b-4b (Mean)	**30.77**	42.19	46.65	51.24	**38.61**	18.61	30.70	**23.31**	23.53	66.50	**88.67**	35.09	2.15	80.00	43.16	46.72	41.74
ExpAttn	23.27	32.95	36.13	40.83	31.71	13.74	27.81	21.46	22.15	65.75	84.11	33.95	5.03	77.22	36.30	57.69	38.13
SnapKV	18.27	21.07	25.80	29.97	17.90	6.88	24.52	18.96	18.22	47.50	85.01	38.58	3.27	46.62	**58.34**	55.63	32.28
SnapKV-QA	28.50	34.45	45.16	**53.20**	34.12	18.53	25.11	22.86	17.56	51.50	85.33	38.70	2.19	**94.20**	**58.34**	58.00	41.73
KVzip	25.80	41.38	49.59	46.31	27.10	15.05	28.73	22.71	23.07	69.50	83.99	37.62	3.79	68.67	49.10	**59.16**	40.72
KVzipFast	20.72	35.80	47.81	45.06	29.11	15.01	27.96	23.00	22.73	66.00	83.63	36.81	3.33	41.96	54.86	56.91	38.17
AM-Fast	28.23	44.29	**49.74**	52.61	36.39	**25.60**	30.70	23.19	**24.02**	73.50	85.27	**39.87**	**5.67**	86.50	54.38	52.05	**44.50**
AM-Slow	25.12	**45.13**	47.57	52.98	36.38	24.04	**30.99**	23.07	23.99	**74.00**	85.51	39.39	4.05	87.50	55.42	52.34	44.22
4x compression																	
LCLM 0.6b-4b (Mean)	**29.73**	**46.84**	47.21	52.87	**42.83**	23.89	**30.93**	23.10	23.99	69.00	**88.25**	34.58	**6.14**	99.00	53.14	57.64	45.57
ExpAttn	27.13	38.72	44.59	49.87	36.76	17.10	29.42	22.63	23.12	71.00	84.55	35.97	3.67	87.29	42.13	57.28	41.95
SnapKV	21.88	30.57	30.90	37.39	26.71	11.58	27.18	20.75	20.71	59.00	84.86	**39.63**	3.00	78.86	**61.53**	56.36	38.18
SnapKV-QA	29.55	39.75	45.91	53.23	37.11	17.92	27.85	23.22	20.13	66.50	84.94	38.60	2.96	93.54	**61.53**	57.14	43.74
KVzip	28.38	45.62	48.61	**54.37**	40.26	18.45	29.86	23.03	23.86	73.00	85.37	39.15	2.28	93.92	59.70	58.92	45.30
KVzipFast	28.61	44.93	48.54	51.89	36.63	18.63	30.24	22.98	23.67	73.00	84.87	38.86	4.55	94.71	60.82	**59.66**	45.16
AM-Fast	26.27	44.45	**49.18**	53.99	37.78	22.66	30.74	22.83	23.89	**74.50**	83.87	39.06	5.50	99.00	58.13	55.85	45.48
AM-Slow	26.30	44.98	47.17	54.16	41.41	**24.10**	30.57	**23.40**	**24.05**	**74.50**	85.19	39.39	3.11	**99.17**	58.30	56.67	**45.78**

Table 14 LongBench Chinese (cn5) per-subtask. Rows grouped by compression ratio (16x, 8x, 4x).

Model	mfq	dure	vcsum	lsht	pr	AVG
Qwen 4B Instruct	61.10	25.32	11.19	45.75	89.33	46.54
16x compression						
LCLM 0.6b-4b (Mean)	39.94	27.00	10.91	23.98	55.09	31.38
ExpAttn	25.52	16.37	11.22	17.75	31.67	20.51
SnapKV	20.36	16.77	10.44	35.00	8.93	18.30
SnapKV-QA	**48.61**	21.46	10.44	**41.50**	**89.28**	**42.26**
KVzip	25.95	16.30	12.50	23.75	7.12	17.12
KVzipFast	30.95	17.04	10.92	19.00	7.47	17.08
AM-Fast	48.55	27.13	13.62	39.50	28.76	31.51
AM-Slow	48.25	**28.02**	**13.75**	39.25	30.00	31.85
8x compression						
LCLM 0.6b-4b (Mean)	42.13	27.05	11.17	30.39	72.50	36.65
ExpAttn	32.75	18.70	11.02	21.00	69.32	30.56
SnapKV	26.07	19.01	10.81	**43.25**	26.78	25.18
SnapKV-QA	58.07	23.37	10.81	42.50	**92.29**	**45.41**
KVzip	58.26	25.30	11.98	40.75	44.17	36.09
KVzipFast	57.35	24.84	12.77	24.00	35.60	30.91
AM-Fast	**59.11**	**28.47**	13.08	42.75	64.21	41.52
AM-Slow	58.88	28.40	**13.18**	42.75	62.34	41.11
4x compression						
LCLM 0.6b-4b (Mean)	44.81	26.49	10.95	29.00	87.29	39.71
ExpAttn	54.32	23.86	11.63	34.25	83.81	41.57
SnapKV	33.48	20.48	11.07	**44.75**	61.32	34.22
SnapKV-QA	60.69	24.33	11.07	43.50	91.34	46.19
KVzip	62.35	26.90	11.86	43.75	92.37	47.45
KVzipFast	**63.62**	27.23	11.95	44.50	96.02	**48.66**
AM-Fast	58.73	28.45	**12.97**	43.75	**98.03**	48.39
AM-Slow	59.04	**29.09**	12.61	43.25	96.37	48.07

F.5 LR Sweep

The tables in this subsection and in Appendix F.6 report the sweep and ablation models of Section 5 and Appendix D, which were trained with an earlier version of our SFT data, so their scores differ from those in Table 9.

F.5.1 Continual Pre-training LR sweep

Table 15 Continual Pre-training LR sweep summary (16 × ratio compression, bidirectional mask, mlp adapter, $O = 0$).

Model	RULER 4k	RULER 8k	RULER 16k	LB en16	LB cn5	LongHealth5	GSM8K
1.0×10^{-6}	62.33	58.75	56.65	33.29	22.88	60.25	66.94
3.0×10^{-6}	70.73	65.41	63.55	32.89	22.09	62.50	72.25
1.0×10^{-5}	**71.62**	**66.83**	**63.63**	**33.52**	**23.94**	**62.75**	**74.53**

Table 16 Continual Pretraining LR sweep – RULER per-task

Model	ns1	ns2	ns3	nm1	nm2	nm3	nmv	nmq	vt	cwe	fwe	qa1	qa2	AVG
Context length 4096														
1.0×10^{-6}	97.40	66.40	31.40	62.00	47.20	17.40	71.70	62.10	84.88	**69.16**	88.00	62.80	49.80	62.33
3.0×10^{-6}	**98.80**	**90.60**	**42.60**	**81.40**	54.40	21.80	84.50	81.15	**91.48**	65.72	88.07	65.20	**53.80**	70.73
1.0×10^{-5}	97.80	83.00	39.40	**81.40**	**64.60**	**29.80**	**85.55**	**81.30**	90.96	68.24	**91.07**	**66.20**	51.80	**71.62**
Context length 8192														
1.0×10^{-6}	96.20	77.80	28.00	60.80	40.00	12.40	60.40	57.95	84.84	**38.70**	**96.40**	**62.20**	48.00	58.75
3.0×10^{-6}	99.00	88.00	40.80	**75.60**	44.20	14.40	**78.55**	**76.15**	**91.40**	34.18	95.47	61.60	**51.00**	65.41
1.0×10^{-5}	**99.20**	**89.80**	**43.60**	73.60	**58.60**	**19.60**	76.35	74.10	91.04	36.74	95.33	61.80	49.00	**66.83**
Context length 16384														
1.0×10^{-6}	98.60	78.00	32.00	64.60	38.40	8.00	60.75	58.50	81.96	**12.66**	**99.53**	56.60	46.80	56.65
3.0×10^{-6}	**99.40**	**90.20**	**46.60**	74.00	38.00	11.20	77.40	**77.15**	**91.24**	11.82	**99.53**	59.00	50.60	63.55
1.0×10^{-5}	98.80	84.60	42.00	**75.60**	**49.80**	**12.80**	**78.40**	74.60	89.12	9.88	99.40	**61.20**	**51.00**	**63.63**

Table 17 Continual Pretraining LR sweep – LongBench English (en16) per-subtask.

Model	narr	qasp	mfq	hpot	2wiki	musq	govr	qmsm	mnews	trec	triv	samsm	pcnt	pr	lcc	rbp	AVG
1.0×10^{-6}	22.98	37.30	**42.23**	37.19	**31.51**	**14.42**	28.61	22.15	22.72	**31.00**	85.35	**33.31**	1.50	**50.00**	35.85	36.53	33.29
3.0×10^{-6}	**23.67**	38.81	41.70	**38.09**	29.49	13.43	**29.25**	**22.98**	**22.98**	30.00	86.44	32.39	0.50	41.50	36.19	38.82	32.89
1.0×10^{-5}	23.32	**40.05**	42.18	35.36	31.18	12.64	29.12	22.95	22.85	29.50	**88.68**	32.28	**4.29**	44.00	**38.07**	**39.92**	**33.52**

Table 18 Continual Pretraining LR sweep – LongBench Chinese (cn5) per-subtask.

Model	mfq	dure	vcsum	lsht	pr	AVG
1.0×10^{-6}	38.79	24.13	**11.00**	**21.00**	19.50	22.88
3.0×10^{-6}	36.39	25.97	10.85	17.75	19.50	22.09
1.0×10^{-5}	**39.65**	**27.07**	10.99	19.50	**22.50**	**23.94**

F.5.2 Post-training LR sweep

Table 19 Post-training LR sweep summary (16 × ratio compression, causal mask, mlp adapter, continual pre-training LR $= 1.0 \times 10^{-5}$).

Model	RULER 4k	RULER 8k	RULER 16k	LB en16	LB cn5	LongHealth5	GSM8K
2.0×10^{-5}	72.99	69.08	63.89	37.98	33.08	66.75	80.89
3.0×10^{-5}	75.06	70.23	65.91	**39.08**	31.74	67.50	**81.05**
5.0×10^{-5}	**75.17**	**71.06**	**66.22**	38.72	**34.16**	**69.00**	80.97

Table 20 Post-training LR sweep – RULER per-task

Model	ns1	ns2	ns3	nm1	nm2	nm3	nmv	nmq	vt	cwe	fwe	qa1	qa2	AVG
Context length 4096														
2.0×10^{-5}	94.60	89.20	40.60	82.40	80.20	46.00	72.20	75.55	83.56	78.94	89.40	**66.40**	**49.80**	72.99
3.0×10^{-5}	**96.40**	89.40	54.60	**85.00**	80.80	48.60	**72.45**	**80.20**	**83.92**	**80.54**	88.27	66.00	49.60	75.06
5.0×10^{-5}	96.00	**90.00**	**58.20**	85.00	**87.00**	**50.60**	71.10	74.65	82.12	79.04	**91.93**	64.20	47.40	**75.17**
Context length 8192														
2.0×10^{-5}	93.60	86.20	51.20	80.20	67.00	32.20	68.00	**77.70**	**80.76**	54.82	**95.73**	**63.40**	**47.20**	69.08
3.0×10^{-5}	**96.00**	**90.20**	55.20	**83.40**	65.40	36.60	**70.25**	77.65	80.52	55.44	95.67	61.80	44.80	70.23
5.0×10^{-5}	94.60	89.80	**57.60**	83.40	**76.20**	**37.40**	69.25	75.15	73.32	**63.56**	95.33	61.20	47.00	**71.06**
Context length 16384														
2.0×10^{-5}	92.20	85.80	47.00	**75.60**	55.60	29.00	68.10	68.25	**79.44**	24.10	98.67	**61.00**	**45.80**	63.89
3.0×10^{-5}	95.80	88.00	**57.40**	75.60	62.20	**30.00**	**70.70**	**73.05**	77.32	24.78	**98.73**	59.40	43.80	65.91
5.0×10^{-5}	**96.40**	**90.60**	53.80	74.60	**72.40**	28.20	67.50	69.70	67.60	**38.38**	98.53	60.20	43.00	**66.22**

Table 21 Post-training LR sweep – LongBench English (en16) per-subtask.

Model	narr	qasp	mfq	hpot	2wiki	musq	govr	qmsm	mnews	trec	triv	samsm	pcnt	pr	lcc	rbp	AVG
2.0×10^{-5}	23.54	40.34	**48.52**	46.17	36.93	21.76	**28.59**	22.21	22.76	**49.00**	87.88	32.31	3.50	74.00	34.65	35.54	37.98
3.0×10^{-5}	**25.57**	39.80	47.98	**47.96**	**38.44**	21.17	28.56	**22.42**	**23.48**	45.00	**89.43**	**32.50**	3.00	81.50	**37.47**	**41.05**	**39.08**
5.0×10^{-5}	23.12	**42.47**	47.64	46.80	37.13	**24.03**	28.27	22.04	22.80	44.00	87.10	32.15	**4.50**	**85.50**	35.45	36.57	38.72

Table 22 Post-training LR sweep – LongBench Chinese (cn5) per-subtask.

Model	mfq	dure	vcsum	lsht	pr	AVG
2.0×10^{-5}	47.53	**24.26**	**10.77**	**20.83**	62.00	33.08
3.0×10^{-5}	46.83	23.63	10.55	18.67	59.00	31.74
5.0×10^{-5}	**47.61**	21.95	10.38	18.85	**72.00**	**34.16**

F.6 Architecture Comparison

F.6.1 Pooling Operator

Table 23 Pooling operator: Mean vs Concat ($W = 1024$, causal, continual pre-training LR = 1.0×10^{-5}). These models were trained with an earlier version of our SFT data, so their scores differ from those in Table 9.

Model	RULER 4k	RULER 8k	RULER 16k	LB en16	LB cn5	LongHealth5	GSM8K
4x compression							
Mean	91.76	91.03	89.96	**46.04**	40.92	76.25	**91.05**
Concat	**92.30**	**91.50**	**90.50**	45.39	**41.80**	**82.50**	89.90
8x compression							
Mean	85.42	84.48	82.47	42.23	34.61	71.75	87.26
Concat	**87.20**	**86.80**	**84.40**	**42.32**	**36.46**	**79.50**	**87.40**
16x compression							
Mean	**75.06**	70.23	**65.91**	**39.08**	**31.74**	67.50	**81.05**
Concat	74.50	**70.30**	64.90	36.24	25.88	**76.30**	78.90

Table 24 Pooling operator: RULER per-task at 4k context.

Model	ns1	ns2	ns3	nm1	nm2	nm3	nmv	nmq	vt	cwe	fwe	qa1	qa2	AVG
4x compression														
Mean	**99.40**	99.60	93.40	99.00	99.40	93.80	98.25	**99.20**	**94.64**	91.42	86.73	**79.20**	58.80	91.76
Concat	99.20	**99.80**	**94.40**	**99.40**	**99.80**	**97.60**	**98.80**	**99.20**	93.90	**92.80**	**87.60**	77.20	**59.80**	**92.27**
8x compression														
Mean	**98.60**	94.20	64.20	93.00	**98.20**	82.20	92.30	90.70	87.40	90.80	**90.60**	73.20	**55.00**	85.42
Concat	98.40	**97.00**	**76.60**	**94.60**	95.60	**86.00**	**93.80**	**92.30**	**90.70**	**92.80**	88.40	**74.80**	53.20	**87.25**
16x compression														
Mean	**96.40**	89.40	54.60	**85.00**	80.80	48.60	**72.45**	**80.20**	**83.92**	**80.54**	88.27	**66.00**	**49.60**	**75.06**
Concat	**96.40**	**91.80**	**62.40**	83.80	**84.20**	**54.00**	59.50	77.50	79.50	74.30	**90.10**	65.40	49.00	74.45

Table 25 Pooling operator: RULER per-task at 8k context.

Model	ns1	ns2	ns3	nm1	nm2	nm3	nmv	nmq	vt	cwe	fwe	qa1	qa2	AVG
4x compression														
Mean	**99.40**	**99.80**	95.60	98.60	98.40	91.60	**98.80**	98.95	94.64	82.04	**90.40**	**76.20**	**59.00**	91.03
Concat	99.20	99.60	**96.60**	**99.20**	**99.60**	**96.00**	98.60	**99.20**	**94.90**	**84.20**	89.80	75.20	57.40	**91.50**
8x compression														
Mean	**98.80**	**96.40**	71.40	90.80	**96.80**	74.60	90.40	90.95	89.40	79.82	**95.67**	70.60	52.60	84.48
Concat	98.60	95.80	**82.00**	**92.00**	95.40	**81.00**	**93.10**	**92.30**	**90.00**	**88.00**	94.60	**71.60**	**53.60**	**86.77**
16x compression														
Mean	96.00	90.20	55.20	**83.40**	65.40	36.60	**70.25**	77.65	**80.52**	**55.44**	**95.67**	61.80	**44.80**	70.23
Concat	**97.60**	**93.20**	**61.80**	82.80	**72.80**	**41.20**	62.10	**78.00**	68.10	54.60	94.70	**62.60**	**44.80**	**70.33**

Table 26 Pooling operator: RULER per-task at 16k context.

Model	ns1	ns2	ns3	nm1	nm2	nm3	nmv	nmq	vt	cwe	fwe	qa1	qa2	AVG
4x compression														
Mean	**99.60**	**99.60**	94.80	99.20	98.20	90.40	98.15	**98.80**	91.40	67.34	**99.53**	**75.60**	56.80	89.96
Concat	99.20	99.20	**95.60**	**99.60**	**99.00**	**92.40**	**98.30**	98.70	**92.80**	**69.90**	98.50	**75.60**	**57.60**	**90.49**
8x compression														
Mean	**99.20**	95.00	72.60	**88.80**	**94.20**	69.20	88.10	90.10	89.04	67.34	97.73	68.40	**52.40**	82.47
Concat	98.80	**95.80**	**82.00**	88.20	90.20	**73.20**	**91.00**	**90.90**	**90.00**	**77.50**	**98.50**	**69.40**	52.20	**84.44**
16x compression														
Mean	95.80	88.00	57.40	**75.60**	62.20	**30.00**	**70.70**	**73.05**	**77.32**	24.78	98.73	59.40	**43.80**	**65.91**
Concat	**96.80**	**90.60**	**65.00**	**75.60**	**65.00**	28.80	63.10	70.30	58.20	**26.60**	**98.90**	**60.80**	43.60	64.87

Table 27 Pooling operator: LongBench English (en16) per-subtask.

Model	narr	qasp	mfq	hpot	2wiki	musq	govr	qmsm	mnews	trec	triv	samsm	pcnt	pr	lcc	rbp	AVG
4x compression																	
Mean	**30.74**	**46.76**	**49.76**	52.90	**42.75**	28.61	**30.12**	**22.89**	23.83	**68.50**	87.90	34.10	6.67	**100.00**	**52.01**	**59.17**	**46.04**
Concat	28.30	46.00	48.70	**56.80**	42.40	**29.70**	29.80	22.80	**24.00**	66.00	**88.50**	**34.20**	**7.50**	98.00	49.20	54.30	45.39
8x compression																	
Mean	28.60	**43.92**	**51.54**	**54.02**	**45.36**	28.05	29.05	22.85	23.56	62.00	86.39	**33.17**	**5.45**	**71.50**	43.57	46.63	42.23
Concat	**30.30**	42.70	50.70	50.80	44.30	**28.20**	**29.20**	**22.90**	**23.90**	**62.50**	**89.00**	33.10	4.60	70.50	**45.20**	**49.20**	**42.32**
16x compression																	
Mean	**25.57**	**39.80**	47.98	47.96	**38.44**	21.17	28.56	**22.42**	**23.48**	45.00	**89.43**	32.50	3.00	**81.50**	37.47	**41.05**	**39.08**
Concat	23.30	36.80	**48.00**	**49.00**	38.30	**22.60**	**28.60**	22.40	23.10	**46.50**	88.20	**32.70**	**3.50**	37.00	**39.90**	39.90	36.24

Table 28 Pooling operator: LongBench Chinese (cn5) per-subtask.

Model	mfq	dure	vcsum	lsht	pr	AVG
4x compression						
Mean	45.52	26.07	**11.00**	25.00	**97.00**	40.92
Concat	**46.20**	**26.80**	**11.00**	**30.00**	95.00	**41.80**
8x compression						
Mean	51.27	**24.83**	10.79	25.65	60.50	34.61
Concat	**53.10**	24.10	**10.90**	**26.70**	**67.50**	**36.46**
16x compression						
Mean	46.83	**23.63**	**10.55**	18.67	**59.00**	**31.74**
Concat	**51.00**	23.50	10.10	**19.30**	25.50	25.88

F.6.2 Adapter & Overlap

Table 29 Adapter & overlap comparison (16 × compression, mean pooling, $W = 1024$, continual pre-training LR $= 1.0 \times 10^{-5}$).

Model	RULER 4k	RULER 8k	RULER 16k	LB en16	LB cn5	LongHealth5	GSM8K
Causal-MLP-O0	**75.06**	**70.23**	**65.91**	**39.08**	**31.74**	**67.50**	**81.05**
Bidirectional-MLP-O0	71.62	66.83	63.63	33.52	23.94	62.75	74.53
Bidirectional-ATTN-MLP-O256	71.91	68.11	64.96	36.28	26.54	64.00	73.77
Bidirectional-MLP-O256	62.03	58.66	58.24	36.97	26.63	63.25	74.98

Table 30 Adapter & overlap comparison – RULER per-task

Model	ns1	ns2	ns3	nm1	nm2	nm3	nmv	nmq	vt	cwe	fwe	qa1	qa2	AVG
Context length 4096														
Causal-MLP-O0	96.40	**89.40**	**54.60**	**85.00**	**80.80**	**48.60**	72.45	80.20	83.92	**80.54**	88.27	66.00	49.60	**75.06**
Bidirectional-MLP-O0	97.80	83.00	39.40	81.40	64.60	29.80	85.55	81.30	90.96	68.24	**91.07**	**66.20**	**51.80**	71.62
Bidirectional-ATTN-MLP-O256	**99.80**	80.40	50.20	78.60	62.20	26.20	**85.70**	**81.35**	95.00	80.00	87.00	61.80	46.60	71.91
Bidirectional-MLP-O256	99.00	60.60	21.60	61.80	55.40	17.60	69.90	61.60	**97.12**	57.66	87.27	66.00	50.80	62.03
Context length 8192														
Causal-MLP-O0	96.00	**90.20**	55.20	**83.40**	**65.40**	**36.60**	70.25	**77.65**	80.52	**55.44**	**95.67**	61.80	44.80	**70.23**
Bidirectional-MLP-O0	99.20	89.80	43.60	73.60	58.60	19.60	76.35	74.10	91.04	36.74	95.33	61.80	**49.00**	66.83
Bidirectional-ATTN-MLP-O256	**100.00**	86.00	**58.60**	72.00	52.60	25.60	**76.75**	74.65	96.12	43.78	91.93	**63.20**	44.20	68.11
Bidirectional-MLP-O256	98.80	67.20	25.40	61.20	47.20	15.60	60.55	54.15	**96.68**	28.76	95.60	62.80	48.60	58.66
Context length 16384														
Causal-MLP-O0	95.80	88.00	**57.40**	75.60	**62.20**	**30.00**	70.70	73.05	77.32	**24.78**	98.73	59.40	43.80	**65.91**
Bidirectional-MLP-O0	98.80	84.60	42.00	75.60	49.80	12.80	**78.40**	74.60	89.12	9.88	99.40	**61.20**	**51.00**	63.63
Bidirectional-ATTN-MLP-O256	**100.00**	**89.00**	57.00	**78.20**	42.40	17.60	78.25	**77.40**	93.80	10.12	**99.47**	56.20	45.00	64.96
Bidirectional-MLP-O256	97.80	75.40	28.60	68.20	36.80	12.80	62.95	60.65	**96.00**	12.48	99.20	59.80	46.40	58.24

Table 31 Adapter & overlap comparison – LongBench English (en16) per-subtask.

Model	narr	qasp	mfq	hpot	2wiki	musq	govr	qmsm	mnews	trec	triv	samsm	pcnt	pr	lcc	rbp	AVG
Causal-MLP-O0	25.57	39.80	**47.98**	47.96	**38.44**	**21.17**	28.56	22.42	**23.48**	**45.00**	**89.43**	**32.50**	3.00	**81.50**	37.47	41.05	**39.08**
Bidirectional-MLP-O0	23.32	**40.05**	42.18	35.36	31.18	12.64	**29.12**	**22.95**	22.85	29.50	88.68	32.28	4.29	44.00	38.07	39.92	33.52
Bidirectional-ATTN-MLP-O256	**27.45**	38.87	42.56	**52.31**	36.98	20.22	27.98	22.44	23.01	27.50	87.23	32.16	**6.00**	51.50	41.22	43.00	36.28
Bidirectional-MLP-O256	23.20	39.67	42.77	46.56	31.94	16.41	28.64	22.50	22.65	40.00	86.67	31.80	4.50	67.50	**41.76**	**44.89**	36.97

Table 32 Adapter & overlap comparison – LongBench Chinese (cn5) per-subtask.

Model	mfq	dure	vcsum	lsht	pr	AVG
Causal-MLP-O0	**46.83**	23.63	10.55	18.67	**59.00**	**31.74**
Bidirectional-MLP-O0	39.65	**27.07**	10.99	19.50	22.50	23.94
Bidirectional-ATTN-MLP-O256	44.35	23.26	10.75	18.83	35.50	26.54
Bidirectional-MLP-O256	38.06	24.02	**11.05**	**22.50**	37.50	26.63

F.6.3 Mask-type Comparison

Table 33 Mask-type ablation (16 × compression, mean pooling, W=1024, continual pre-training LR $= 1.0 \times 10^{-5}$).

Mask	RULER 4k	RULER 8k	RULER 16k	LB en16	LB cn5	LongHealth5	GSM8K
Bidirectional	71.62	66.83	63.63	33.52	23.94	62.75	74.53
Causal	**75.06**	**70.23**	**65.91**	**39.08**	**31.74**	**67.50**	**81.05**

F.6.4 Window-size Comparison

Table 34 Window-size ablation (16 × compression, mean pooling, continual pre-training LR $= 1.0 \times 10^{-6}$). W=16 and W=256 use causal attention; W=1024 uses bidirectional attention. We do not run W=1024 with causal attention at LR $= 1.0 \times 10^{-6}$. However, Table 33 shows that causal masking is strictly better than bidirectional masking, and W=1024 (bidirectional) already outperforms W=256 (causal) on RULER tasks while also achieving lower pre-training loss. We therefore conclude that W=1024 is the better setting.

Model	RULER 4k	RULER 8k	RULER 16k	LB en16	LB cn5	LongHealth5	GSM8K
Mean W16 (causal)	24.85	23.87	22.54	29.62	22.62	56.50	46.55
Mean W256 (causal)	58.37	54.39	50.18	**34.15**	**25.97**	**62.25**	**76.80**
Mean W1024 (bidir)	**62.33**	**58.75**	**56.65**	33.29	22.88	60.25	66.94

F.6.5 Encoder-representation Comparison

Table 35 Encoder-representation ablation (16 $\times$ compression, EOS pooling, causal, continual pre-training LR $= 1.0 \times 10^{-6}$).

Encoder init	RULER 4k	RULER 8k	RULER 16k	LB en16	LB cn5	LongHealth5	GSM8K
LLM (Instruct)	38.29	34.39	33.50	30.48	22.73	**58.25**	54.51
Embed (Qwen3-Emb)	**42.00**	**38.76**	**38.12**	**31.72**	**25.32**	58.00	**59.36**

F.7 NIAH Agent Results

Table 36 RULER NIAH (retrieval) per-task results for our LCLM 0.6B-4B model at 16x compression, with and without the agentic expand-or-answer loop, across context lengths. Δ *agent* is the cell-wise gain over the non-agentic baseline. Each cell averages 50 questions per task and context length, matching Figure 6.

Model	ns1	ns2	ns3	nm1	nm2	nm3	nmv	nmq	AVG
					4k context				
LCLM 0.6B-4B	100.00	88.00	70.00	90.00	86.00	56.00	72.00	85.50	80.94
LCLM 0.6B-4B agent	**100.00**	**100.00**	**100.00**	**100.00**	**100.00**	**100.00**	**95.50**	**100.00**	**99.44**
Δ *agent*	*+0.00*	*+12.00*	*+30.00*	*+10.00*	*+14.00*	*+44.00*	*+23.50*	*+14.50*	*+18.50*
					8k context				
LCLM 0.6B-4B	96.00	90.00	58.00	82.00	88.00	44.00	70.00	79.00	75.88
LCLM 0.6B-4B agent	**100.00**	**100.00**	**100.00**	**100.00**	**98.00**	**100.00**	**95.00**	**99.50**	**99.06**
Δ *agent*	*+4.00*	*+10.00*	*+42.00*	*+18.00*	*+10.00*	*+56.00*	*+25.00*	*+20.50*	*+23.19*
					16k context				
LCLM 0.6B-4B	92.00	82.00	60.00	86.00	60.00	26.00	68.00	82.50	69.56
LCLM 0.6B-4B agent	**100.00**	**100.00**	**100.00**	**98.00**	**86.00**	**86.00**	**93.00**	**97.50**	**95.06**
Δ *agent*	*+8.00*	*+18.00*	*+40.00*	*+12.00*	*+26.00*	*+60.00*	*+25.00*	*+15.00*	*+25.50*