
LoRC: Low-Rank Compression for LLMs KV Cache with a Progressive Compression Strategy

Anonymous Author(s)

Affiliation

Address

email

Abstract

1 The Key-Value (KV) cache is a crucial component in serving transformer-based
2 autoregressive large language models (LLMs), enabling faster inference by stor-
3 ing previously computed KV vectors. However, its memory consumption scales
4 linearly with sequence length and batch size, posing a significant bottleneck in
5 LLM deployment. Existing approaches to mitigate this issue include: (1) efficient
6 attention variants integrated in upcycling stages, which requires extensive parame-
7 ter tuning thus unsuitable to pre-trained LLMs; (2) KV cache compression at test
8 time, primarily through token eviction policies, which often overlook inter-layer
9 dependencies and can be task-specific.

10 This paper introduces an orthogonal approach to KV cache compression. We
11 propose a *low-rank approximation* of KV weight matrices, allowing for plug-
12 in integration with existing transformer-based LLMs without model retraining.
13 To effectively compress KV cache at the weight level, we adjust for layerwise
14 sensitivity and introduce a *progressive compression* strategy, which is supported by
15 our theoretical analysis on how compression errors accumulate in deep networks.
16 Our method is designed to function without model tuning in upcycling stages or
17 task-specific profiling in test stages. Extensive experiments with LLaMA models
18 ranging from 8B to 70B parameters across various tasks show that our approach
19 significantly reduces the GPU memory footprint while maintaining performance.

20 1 Introduction

21 Autoregressive large language models (LLMs) such as GPT (Achiam et al., 2023), PaLM (Chowdhery
22 et al., 2023), and LLaMA (Touvron et al., 2023), built upon transformer architectures (Vaswani
23 et al., 2017), have shown remarkable capabilities across a wide range of tasks. However, the
24 attention mechanism underpinning those models poses significant challenges to the efficiency of
25 their deployment, particularly the management of the Key-Value (KV) cache. The KV cache is
26 originally designed to accelerate the generation process by storing intermediate attention KV vectors,
27 thus avoiding recomputation of shared prefixes for each autoregressively generated token. Despite
28 reducing computational overhead, the KV cache significantly increases memory footprints, as its
29 size scales linearly with both sequence length and batch size. This drives the need for KV cache
30 compression to enable cost-effective deployment of LLMs across various devices and platforms.

31 To address the overhead of the original attention mechanism, one prominent line of work aims
32 to design more efficient attention variants, such as multi-query attention (MQA) (Shazeer, 2019)
33 and group-query attention (GQA) (Ainslie et al., 2023), which inherently reduce the corresponding
34 KV cache. Nevertheless, those techniques typically require upcycling existing models. Without
35 proper training, their direct application often results in degraded performance (Ribar et al., 2023;
36 Ainslie et al., 2023; Liu et al., 2024b), thereby making them unsuitable for deployment in resource-

37 constrained environments. Recently, Liu et al. (2024a) design a multi-head latent attention (MLA) for
38 efficient inference, utilizing low-rank key-value union compression to reduce KV cache. However,
39 similar to MQA and GQA, MLA is also integrated during the model’s training cycle, thus not directly
40 applicable to pre-trained LLMs.

41 In contrast, another line of work focuses on KV cache compression at test time, primarily achieved
42 by dropping tokens while leaving the backbone model intact. Several works design the token eviction
43 policy based on accumulated attention scores(Sheng et al., 2023; Zhang et al., 2024b; Liu et al.,
44 2024b), or heuristics such as special tokens or and relative distance between tokens (Ge et al., 2023)
45 However, these methods either ignore the inter-layer dependency or require attention pattern analysis,
46 and the resulting eviction policy can be task-specific.

47 In this paper, we propose to compress KV cache from an orthogonal perspective, *i.e.*, the KV weight
48 matrices. As the KV weight matrices are typically characterized by low-rank properties, we perform
49 a *low-rank approximation* to reduce their dimension and thus compress the resulting KV cache.
50 Recognizing that compressed KV caches inevitably introduce information loss to subsequent layers,
51 and that sensitivity to input changes varies across layers, we introduce a *progressive* compression
52 strategy. This approach is grounded in the calculation of cumulative condition numbers for KV weight
53 matrices across different layers, reflecting their sensitivity and guiding the compression strategy.
54 Theoretically, we derive error bounds for both individual layer compression and error propagation
55 through the network. These theoretical results reveal that errors introduced in earlier (shallower)
56 layers are amplified more significantly than those in deeper layers, and informs our progressive
57 compression strategy.

58 Our method is designed for straightforward implementation, requiring neither model profiling nor
59 detailed inspection of the attention structure. It can be directly applied to pre-trained LLMs by
60 extracting weight matrices and leveraging their inherent properties to swiftly determine optimal
61 layer-wise compression. This approach offers a practical and efficient solution for enhancing LLM
62 performance in memory-constrained deployment scenarios, without the need for model retraining or
63 complex eviction strategy composition.

64 We evaluate our method on 8B, 13B, and 70B LLaMA models that built upon multi-query attention
65 or group-query attention. Experiments across tasks such as commonsense reasoning, reading compre-
66 hension, text summarization, and mathematical reasoning, demonstrate that our approach can reduce
67 substantial GPU memory footprint while maintaining minimal impact on performance.

68 **2 Related Works**

69 **2.1 Attention Mechanism**

70 Attention mechanisms in Transformer models have evolved to enhance efficiency and effectiveness
71 (Vaswani et al., 2017). Multi-Query Attention (MQA)(Shazeer, 2019) reduces memory requirements
72 during decoding, while Grouped-Query Attention (GQA) (Ainslie et al., 2023) balances efficiency
73 and performance by sharing key and value heads among query groups. Recently, Liu et al. (2024a)
74 introduced Multi-head Latent Attention (MLA), using low-rank key-value union compression to
75 optimize inference. However, these approaches are typically integrated during model training, limiting
76 their applicability to pre-trained LLMs. Parallel research efforts have targeted inference efficiency
77 improvements. For example, Pope et al. (2023) developed multi-dimensional partitioning techniques,
78 and de Jong et al. (2022) optimized the Fusion-in-Decoder (FiD) approach (Izacard & Grave, 2020)
79 for more efficient inference. Holmes et al. (2024) introduces SplitFuse which leverages dynamic
80 prompt and generation decomposition and unification to further improve continuous batching and
81 system throughput. In this paper, we contribute to this line of research by improving inference
82 efficiency through the compression of KV cache. Our approach leverages the low-rank property of
83 the attention weight matrices, offering a plug-and-play method to reduce the memory footprint of
84 LLMs during inference without requiring model retraining.

85 **2.2 KV Cache Compression**

86 As Large Language Models (LLMs) continue to grow in size and complexity, efficient management
87 of their memory usage during inference has become a critical challenge. Early efforts to compress
88 token hidden states (Guan et al., 2022; Sun et al., 2022; Zhou et al., 2020) are limited to non-

89 autoregressive models and require retraining, thus motivating research into pruning tokens in the
 90 KV cache of auto-regressive LLMs. For instance, Mu et al. (2024) learns to compress prompts into
 91 a few special tokens to reduce memory pressure during caching, but this token prediction requires
 92 model retraining and could be an expensive overhead during inference. Several methods design token
 93 eviction policies based on accumulated attention scores (Sheng et al., 2023; Zhang et al., 2024b; Liu
 94 et al., 2024b), or heuristics such as special tokens and relative distance between tokens (Ge et al.,
 95 2023). However, these approaches often overlook inter-layer dependencies, potentially resulting in
 96 task-specific eviction policies that may not generalize well across different applications. In contrast
 97 to token-dropping methods, our study takes a different tack. We focus on compressing the KV cache
 98 from the perspective of weight matrix dimension reduction. Importantly, our progressive compression
 99 strategy carefully addresses the issue of error propagation across compressed layers, a consideration
 100 often ignored in previous methods.

101 A few studies have explored customized cache budgets across different layers in the context of token
 102 dropping, yet no definitive consensus has been reached on the most effective strategies. Zhang
 103 et al. (2024a) suggest increasing compression intensity in higher layers based on the assumption that
 104 these layers contain less critical information. Conversely, Liu et al. (2024b) argue that significant
 105 tokens exhibit greater variability at higher layers, thus larger caches are required to reduce cache
 106 misses. While these approaches demonstrate understanding of layer-specific requirements, they
 107 depend heavily on task-specific attention patterns. Our approach diverges fundamentally by adopting
 108 an orthogonal perspective to compression, focusing on weight matrix dimension reduction rather than
 109 token eviction. This approach enables us to establish error propagation bounds across the network and
 110 to guide our progressive compression strategy effectively. It eliminates the need to analyze attention
 111 patterns for eviction policy design, simplifying implementation and enhancing general applicability
 112 across different LLMs.

113 Concurrently, Liu et al. (2024a) and Yu et al. (2024) modify attention mechanisms to manage KV
 114 caches more efficiently during inference. While these methods align with our philosophy of altering
 115 attention dynamics, they require either pretraining adjustments or extensive model finetuning to
 116 accommodate the modified attention schemas, limiting their practicality in deployed systems. In
 117 contrast, our method requires no such training or fine-tuning, offering a plug-and-play solution that
 118 seamlessly integrates with pre-trained models to deliver efficient compression without compromising
 119 the model’s integrity or performance.

120 3 Preliminary: Attention Mechanism and KV Cache

121 Transformer-based language models use self-attention to weigh the importance of different tokens,
 122 thus allowing for the model to focus on different parts of the input sequence. Given an input
 123 $X \in \mathbb{R}^{N \times D}$, where N is the sequence length and D is the dimensionality of each token’s embedding,
 124 we compute the Query (Q), Key (K), and Value (V) matrices by multiplying X with their respective
 125 weight matrices: $Q = XW_q, K = XW_k, V = XW_v$.

126 Then the attention mechanism is as follows:

$$\text{Attention}(Q, K, V) = \text{softmax} \left(\frac{QK^\top}{\sqrt{d_k}} \right) V. \quad (1)$$

127 Multi-head attention allows the model to jointly attend to information from different representation
 128 subspaces at different positions

$$\text{MultiHead}(Q, K, V) = \text{Concat}(\text{head}^1, \dots, \text{head}^h)W_o, \quad (2)$$

129 where

$$\text{head}^i = \text{Attention}(X(W_q^i)^T, X(W_k^i)^T, X(W_v^i)^T).^1 \quad (3)$$

130 Here, $W_q^i, W_k^i,$ and W_v^i are the weight matrices for the i -th attention head, and W_o is the weight
 131 matrix for the output linear transformation.

132 In autoregressive transformers, the computation of attention scales quadratically (*i.e.*, $\mathcal{O}(N^2)$) with
 133 the sequence length N , as every token in the sequence computes interactions with every other token.

¹This formulation with transposed weight matrices aligns with the implementation found in the models examined in our study. Mathematically, this is equivalent to the standard formulation without transpose. The choice of which form to use depends on implementation details and computational optimizations.

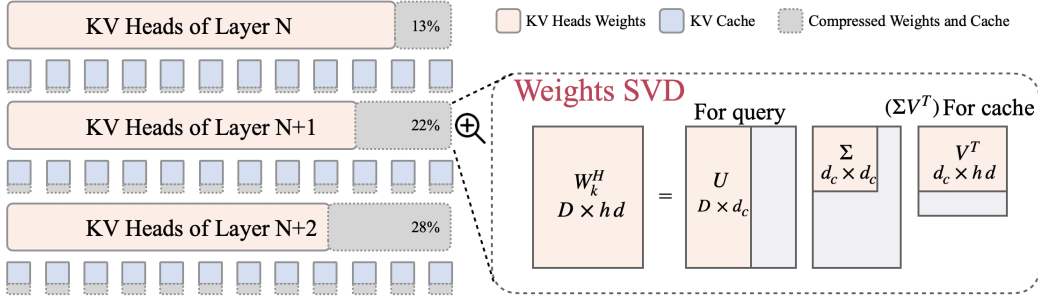


Figure 1: LORC compresses KV-cache by decomposing the KV weight matrices in attention heads. The progressive compression strategy retains more dimension for KV weights in shallow layers and compresses the KV weights in deep layers more aggressively.

134 Such scaling is impractical for very large inputs or real-time applications, where speed and efficiency
 135 are crucial.

136 To address this computational bottleneck, KV caches store the results of previous computations of
 137 the KV matrices. When processing subsequent tokens, the model can retrieve keys and values from
 138 the cache rather than recomputing them, thereby reducing the number of operations to a linear scale
 139 with respect to the sequence length. This method trades off increased memory usage for a reduction
 140 in computational overhead. The size of KV cache per layer is defined as below:

$$C_{k,v} = b \times N \times h \times d, \quad (4)$$

141 where b is the batch size, N is the max sequence length in the batch, h is the number of K/V head
 142 and d is the head dimension. This linear relationship between cache size and sequence length, as well
 143 as batch size, underscores the critical need for efficient compression methods. As described, existing
 144 works that can reduce KV cache consumption either require expensive model training in upcycling
 145 stages or empirical token eviction policy design at test time. In the following section, we present a
 146 novel method for KV cache compression from the perspective of low-rank weight approximation.

147 4 Method

148 We structure this section as follows. In Section 4.1, we detail the process of compressing the KV
 149 cache for a single layer using Singular Value Decomposition (SVD) on weight matrices. Section 4.2
 150 introduces our progressive compression strategy, which determines adaptive compression dimensions
 151 for each layer. Finally, Section 4.3 covers additional considerations for handling various attention
 152 mechanisms, and Section C addresses the implementation details specific to the rotary position
 153 embedding. Figure 1 presents an overview of our method, illustrating the low-rank approximation of
 154 the weight matrix and the progressive compression strategy across layers.

155 4.1 KV Cache Compression via Low-rank Approximation of Weight Matrices

156 Unlike previous approaches that focus on token-level eviction strategies or require model retraining,
 157 we propose a novel method that operates at the weight matrix level in the attention mechanism. This
 158 approach leverages the inherent low-rank properties of these matrices (as shown in Appendix B),
 159 allowing for significant compression without the need for complex token selection algorithms or
 160 time-consuming model tuning. By applying a low-rank approximation to the weight matrices, we
 161 effectively reduce the dimensionality of the KV cache while preserving the essential information flow
 162 through the network.

163 **Key Matrix Compression:** Figure 1 presents how we implement SVD on the key weight matrices.
 164 Specifically, for the i -th head in the MHA attention, we decompose its key matrix $W_k^i \in \mathbb{R}^{D \times d}$ to:

$$\text{SVD}(W_k^i)_{D \times d} = U_{D \times d_c} \Sigma_{d_c \times d_c} V_{d_c \times d}^T = U_{D \times d_c} (\Sigma V^T)_{d_c \times d}. \quad (5)$$

165 For MHA, there are h attention heads, then the decomposition becomes:

$$\text{SVD}(W_k^H)_{D \times hd} = U_{D \times d_c} (\Sigma V^T)_{d_c \times hd} = U_{D \times d_c} \left[(A^1)_{d_c \times d} \quad (A^2)_{d_c \times d} \quad \cdots \quad (A^h)_{d_c \times d} \right], \quad (6)$$

166 where $(A^i)_{d_c \times d}$ is the i -th block in the matrix $(\Sigma V^T)_{d_c \times hd}$.

167 Now we have decomposed the key matrix W_k^i to the multiplication of $U_{D \times d_c}$ and $(\Sigma V^T)_{d_c \times hd}$. We
 168 will multiply X with $(\Sigma V^T)_{hd \times d_c}^T$ as the compressed key, which is stored in the KV cache. Through
 169 this implementation, we effectively update the size of key cache from hd to d_c , where d_c is smaller
 170 than hd , reducing the memory footprint while keeping the essential information intact.

171 For $U_{D \times d_c}$, we incorporate it to the query calculation by updating the original query matrix $W_q^H \in$
 172 $\mathbb{R}^{D \times hd}$ as follows:

$$W_{q'}^H = (W_q^H)_{D \times hd} U_{D \times d_c}. \quad (7)$$

173 Note that the embedding dimension D is equal to the product of the number of attention heads h and
 174 the dimension per head d , i.e., $D = hd$. Consequently, the updated query matrix $W_{q'}^H \in \mathbb{R}^{D \times d_c}$.

175 **Value Matrix Compression:** The decomposition for the value matrix follows a similar structure
 176 to that of the key matrix, with the difference that we integrate its left singular vectors to the output
 177 matrix W_o . Specifically, the value matrix is decomposed as:

$$\text{SVD}(W_v^H)_{D \times hd} = U_{D \times d_c} (\Sigma V^T)_{d_c \times hd} = U_{D \times d_c} \left[(B^1)_{d_c \times d} \quad (B^2)_{d_c \times d} \quad \cdots \quad (B^h)_{d_c \times d} \right] \quad (8)$$

178 where $(B^i)_{d_c \times d}$ is the i -th block in the matrix $(\Sigma V^T)_{d_c \times hd}$. After multiplication with X , the
 179 dimension of the value cache shrinks from hd to d_c , thus reducing memory consumption.

180 In contrast to the key matrix operation, we incorporate $U_{D \times d_c}$ to the output matrix. To achieve this,
 181 we update the output matrix $W_o \in \mathbb{R}^{D \times D}$ as follows:

$$W_{o'} = (U^\top)_{d_c \times D} (W_o)_{D \times D}, \quad (9)$$

182 resulting in an updated output matrix $W_{o'} \in \mathbb{R}^{d_c \times D}$.

183 **Compression Ratio:** The compression strategy effectively reduces the dimensions from $N \times d \times h$
 184 for both keys and values to $N \times d_c$, ensuring data integrity and minimizing overhead. This results in
 185 a layer compression ratio $\rho = \frac{d_c}{h \times d}$, which quantifies the extent of the reduction.

186 4.2 Progressive Compression Strategy

187 Algorithm 1 LORC Algorithm

Require: Pre-trained LLM with L layers

```

1: Initialize cumulative condition numbers  $\tilde{\kappa}_l$ 
2: for  $l = L$  to 1 do
3:   Compute  $\kappa(W_k^l)$  and  $\kappa(W_v^l)$ 
4:    $\tilde{\kappa}_l \leftarrow \prod_{j=l}^L \kappa(W_k^j) \cdot \kappa(W_v^j)$ 
5: end for
6: for  $l = 1$  to  $L$  do
7:    $d_c^l \leftarrow$  Calculate by Eq. 13
8:   if  $\tilde{\kappa}_l >$  threshold then
9:     Skip compression for layer  $l$ 
10:    continue
11:  end if
12:  Key Matrix Compression:
13:    Perform SVD:  $W_k^l = U_k \Sigma_k (V_k^T)$ 
14:     $\tilde{W}_k^l \leftarrow U_k[:, : d_c^l] (\Sigma_k V_k^T)[:, d_c^l, :]$ 
15:     $W_{q'}^l \leftarrow \tilde{W}_k^l U_q[:, : d_c^l]$ 
16:  Value Matrix Compression:
17:    Perform SVD:  $W_v^l = U_v \Sigma_v (V_v^T)$ 
18:     $\tilde{W}_v^l \leftarrow U_v[:, : d_c^l] (\Sigma_v V_v^T)[:, d_c^l, :]$ 
19:     $W_{o'}^l \leftarrow U_v[:, : d_c^l]^T W_o^l$ 
20:  Update KV cache size for layer  $l$ 
21: end for

```

Having established low-rank approximation for compressing weight matrices, we now address its dynamic application across network layers. This approach is necessary due to the varying sensitivity of different layers, which significantly affects overall model efficacy and efficiency.

To tackle this challenge, we propose a *progressive* compression strategy for our low-rank approximation of KV weight matrices. Our intuition is that the compressed shallow layers could lead to cascading errors that propagate and amplify through the network. Therefore, we measure the layer sensitivity by the condition numbers of KV matrices to determine *layer-wise* compression dimensions. This approach accounts for each layer’s sensitivity to perturbations caused by previously compressed layers, ensuring output variations remain within acceptable ranges. This progressive nature allows for more conservative compression in shallow layers and more aggressive compression in deeper layers, minimizing the risk of error accumulation throughout the network. By carefully balancing compression across layers, we maintain model integrity while achieving significant memory savings.

188 **Condition Number and Sensitivity Analysis** To ensure that the change in the output $\mathbf{b}_l = \mathbf{A}_l \mathbf{x}_l$
 189 remains within a specified range when the input \mathbf{x}_l changes due to compression in previous layers, we
 190 need to consider the sensitivity of the output to such changes. Given a weight matrix \mathbf{A}_l , its condition
 191 number plays a crucial role in determining the allowable change in \mathbf{x}_l . The condition number $\kappa(\mathbf{A}_l)$
 192 is defined as:

$$\kappa(\mathbf{A}_l) = \|\mathbf{A}_l\|_2 \cdot \|\mathbf{A}_l^{-1}\|_2 = \frac{\sigma_{\max}(\mathbf{A}_l)}{\sigma_{\min}(\mathbf{A}_l)}, \quad (10)$$

193 where $\sigma_{\max}(\mathbf{A}_l)$ and $\sigma_{\min}(\mathbf{A}_l)$ are the largest and smallest singular values of \mathbf{A}_l , respectively. To
 194 keep the relative change in the output \mathbf{b}_l within a tolerance ϵ , we utilize the standard definition of the
 195 condition number to relate it to the allowable relative change in the input \mathbf{x}_l :

$$\frac{\|\Delta \mathbf{b}_l\|_2}{\|\mathbf{b}_l\|_2} \leq \kappa(\mathbf{A}_l) \cdot \frac{\|\Delta \mathbf{x}_l\|_2}{\|\mathbf{x}_l\|_2} \leq \epsilon. \quad (11)$$

196 Solving for the allowable relative change in \mathbf{x}_l , we obtain: $\frac{\|\Delta \mathbf{x}_l\|_2}{\|\mathbf{x}_l\|_2} \leq \frac{\epsilon}{\kappa(\mathbf{A}_l)}$. This inequality indicates
 197 that the acceptable change in the input \mathbf{x}_l is *inversely proportional* to the condition number $\kappa(\mathbf{A}_l)$
 198 of the layer’s weight matrix. Layers with higher condition numbers are more sensitive to input
 199 perturbations, requiring smaller changes in \mathbf{x}_l to maintain the output within the desired range. Given
 200 the multi-layer structure of transformers, it is essential to consider not just the condition number of a
 201 single layer but the cumulative effect of condition numbers from all preceding layers. This cumulative
 202 measure gives a more holistic view of how perturbations might propagate and amplify as data passes
 203 through successive layers.

204 **Cumulative Condition Number:** To effectively manage this across the network, we calculate the
 205 cumulative condition number as an estimated layer sensitivity, which we then use to derive the
 206 compression dimension. For a model with L layers, we calculate the cumulative condition number
 207 for each layer l by multiplying the condition numbers of the current layer and all subsequent layers:

$$\tilde{\kappa}_l = \prod_{j=l}^L \kappa(W_k^j) \cdot \kappa(W_v^j), \quad (12)$$

208 where W_k^j and W_v^j denote the key and value weight matrices of the j -th layer, respectively. This
 209 cumulative condition number $\tilde{\kappa}_l$ reflects the total amplification of input perturbations from current
 210 layer to the final output layer, encompassing the effects of layers from l to L .

211 **Compression Dimension:** Based on the cumulative condition number, we then adjust the compression
 212 dimensions for each layer to balance the fidelity and compression rate. More sensitive layers
 213 (those with higher cumulative condition numbers) will have less aggressive compression to preserve
 214 information, whereas layers with lower sensitivity can be compressed more substantially without
 215 significantly affecting the overall network performance. We compute the compressed dimension d_c^l
 216 for each layer by scaling $\tilde{\kappa}_l$ using the following function:

$$d_c^l = d_{\max} \times \left[1 - \left(\frac{\max_{i \in [1:L]} \log(\tilde{\kappa}_i) - \log(\tilde{\kappa}_l)}{\max_{i \in [1:L]} \log(\tilde{\kappa}_i) - \min_{i \in [1:L]} \log(\tilde{\kappa}_i)} \right) \times \left(1 - \frac{d_{\min}}{d_{\max}} \right) \right], \quad (13)$$

217 where d_{\max} is the maximum allowable compressed dimension, and d_{\min} is the minimum one.
 218 The logarithmic scale mitigates the effect of large variations in the cumulative condition numbers,
 219 providing a more balanced sensitivity metric across layers. This equation ensures that layers with
 220 higher sensitivity (larger $\tilde{\kappa}_l$) retain more dimensions (larger d_l), while less sensitive layers can be
 221 compressed more aggressively.

222 4.3 Multi-head Attention and Group-query Attention

223 The above derivation in Section 4.1 holds for standard MHA, where the model dimension D equals
 224 to the multiplication of number of head and head dimension $h \times d$. For GQA, the number of KV
 225 heads is reduced as shown in Table 3. To adapt such implementation, we can still follow the above
 226 procedure for cache compression. After fetching the key and value from cache, we just need to repeat
 227 them according to the number of the total attention heads.

228 **5 Error Bounds for KV Cache Compression**

229 In this section, we derive error bounds for our KV cache compression method, considering both
 230 individual layer errors and their propagation through a deep network. These theoretical results provide
 231 insights into how the matrix decomposition-based compression affects the network’s performance
 232 and guide the progressive compression strategy to balance model efficiency and performance.

233 **5.1 Error Bound for Key/Value Matrix Approximation**

234 **Theorem 1** *Let $W \in \mathbb{R}^{m \times n}$ be a weight matrix (either key or value), and let $\tilde{W} \in \mathbb{R}^{m \times n}$ be its*
 235 *rank- k approximation obtained via truncated singular value decomposition (SVD). For any input*
 236 *vector $x \in \mathbb{R}^n$, the error introduced by the approximation is bounded by:*

$$\|Wx - \tilde{W}x\|_2 \leq \sigma_{k+1} \|x\|_2, \quad (14)$$

237 where σ_{k+1} is the $(k + 1)$ -th singular value of W .

238 The proof is provided in Appendix A.1. This theorem quantifies the error introduced at a single layer
 239 due to compressing the weight matrix. The bound indicates that the error is directly proportional
 240 to the $(k + 1)$ -th singular value of W and the norm of the input vector x . Larger singular values
 241 correspond to directions of significant variance in the data, so truncating smaller singular values
 242 (which represent less significant features) minimizes the error introduced by compression.

243 **5.2 Single Layer Error Bound Including Nonlinearities**

244 We now extend the analysis to include the effect of nonlinearities within a single layer. We derive an
 245 error bound that accounts for both the approximation of the weight matrix and the layer’s nonlinear
 246 activation function. For simplicity, we analyze the error introduced by compressing each weight
 247 matrix (key or value) individually.

248 **Theorem 2** *Consider a single layer applying a linear transformation W followed by a nonlinearity*
 249 *ϕ with Lipschitz constant L_ϕ . Let \tilde{W} be the compressed version of W obtained via truncated SVD*
 250 *with rank k . For any input vector $x \in \mathbb{R}^n$, the error at the output of the layer is bounded by:*

$$\left| \phi(Wx) - \phi(\tilde{W}x) \right| \leq L_\phi \sigma_{k+1} \|x\|_2. \quad (15)$$

251 The proof is straightforward by using Theorem 1 and the Lipschitz property of ϕ , we present it as the
 252 base case in the proof of Theorem 3, which is detailed in Appendix A.2.

253 This theorem shows that the error introduced by the compressed weight matrix propagates through
 254 the nonlinearity, scaled by the Lipschitz constant of the activation function. While considering
 255 both matrices simultaneously complicates the bounds due to their interactions within the attention
 256 mechanism, it is still feasible to derive combined error bounds because the attention mechanism
 257 allows us to mathematically bound these interactions. The total error due to simultaneous compression
 258 can be bounded by the sum of their individual approximation errors, scaled by a constant. However,
 259 for simplicity and clarity in the following derivation, we use the simplified version that considers
 260 each matrix individually.

261 **5.3 Error Propagation Bound**

262 **Theorem 3** *Consider an L -layer network where each layer i applies a linear transformation W_i*
 263 *followed by a nonlinearity ϕ with Lipschitz constant L_ϕ . Let \tilde{W}_i be the compressed version of W_i*
 264 *obtained via truncated SVD with rank k_i . The error at the output of the network is bounded by:*

$$\|x_L - \tilde{x}_L\|_2 \leq \sum_{i=1}^L \left(\sigma_{k_i+1}^{(i)} L_\phi^{L-i} \prod_{j=i+1}^L \|W_j\|_2 \right), \quad (16)$$

265 where x_L and \tilde{x}_L are the outputs of the original and compressed networks, respectively; $\sigma_{k_i+1}^{(i)}$ is the
266 $(k_i + 1)$ -th singular value of W_i ; $\|W_j\|_2$ denotes the spectral norm of W_j ; and L_ϕ is the Lipschitz
267 constant of the activation function ϕ .

268 We detail the proof in Appendix A.2. Until now, we have established an upper bound on the
269 cumulative error at the network’s output due to compression of weight matrices across multiple
270 layers. It is important to note that the nonlinearities characterized by the Lipschitz constant L_ϕ
271 represent a simplification. In practice, transformer models like LLaMA incorporate complex nonlinear
272 components, so the exact error propagation may deviate from this simplified bound due to intricate
273 nonlinearities. Despite these complexities, the theorem still offers insights into how compression
274 errors may accumulate in deep networks. Specifically, it reveals that errors introduced in earlier
275 (shallower) layers are amplified more significantly than those in deeper layers because they pass
276 through more subsequent transformations and nonlinearities.

277 This understanding supports our design of a progressive compression strategy, where we compress
278 shallow layers less aggressively than deeper ones. By preserving more information in the early
279 layers (i.e., retaining more singular values), we minimize the initial errors that could be significantly
280 amplified throughout the network. This approach helps maintain overall model performance while
281 still achieving substantial compression in deeper layers, where the impact on the final output is less
282 pronounced due to reduced error amplification.

283 6 Experiment

284 6.1 Models

285 We conduct experiments using two attention mechanisms, Multi-Head Attention (MHA) (Vaswani
286 et al., 2017) and Graph Query Attention (GQA) (Ainslie et al., 2023), across three models: LLaMA-
287 2-13B, LLaMA-3-Instruct-8B, and LLaMA-3-Instruct-70B. The LLaMA-2 family incorporates the
288 MHA mechanism, while the LLaMA-3 family is based on the GQA framework. We list the model
289 specifications in Table 3. Note that for the models based on MHA, the number of KV heads is equal
290 to the number of attention heads, so the weight matrices of KV are square matrices. The models
291 based on GQA use an intermediate number of key-value heads to group the query heads, with an
292 adjustment on the shape of KV weight matrices.

293 6.2 Implementation Details

294 In practice, we set thresholds to exclude compression on layers with high cumulative condition
295 numbers: 30 for LLaMA-3-Instruct-8B, and 90 for LLaMA-2-13B and LLaMA-3-Instruct-70B. The
296 d_{max} equals to the original head dimension, while d_{min} varies based on the target compression ratio.
297 For baseline methods, we have the same refrained layers while applying the uniform compression
298 ratios across compressed layers instead of using a progressive compression strategy.

299 6.3 Dataset

300 We follow Touvron et al. (2023) to evaluate our methods on the following tasks: **BoolQ** (Clark et al.,
301 2019) for reading comprehension, **XSum** (Narayan et al., 2018) for text summarization, **Openbook**
302 **QA** (Mihaylov et al., 2018) for commonsense reasoning, and **GSM8K** (Cobbe et al., 2021) for
303 mathematical reasoning. We use ROUGE score (Lin, 2004) as the evaluation metric for XSum and
304 accuracy for the other tasks. We report 2-shot results for LLaMA-2 models on BoolQ, and 0-shot
305 results for other settings.

306 6.4 Main Results

307 Figure 2 presents our main results on four datasets with different KV cache budgets. Compared to
308 the full-cache model, LORC achieves on-par performance with a significant compression ratio, and
309 the performance degradation is still nearly negligible with a 60% compression ratio on most datasets.
310 When slightly compressed, LORC could even enhance model performance in some cases. Note that
311 our method requires no model training or model profiling, the only efforts are SVD on weight matrices
312 which requires minimal computational cost compared to the LLM inference. Such plug-and-play

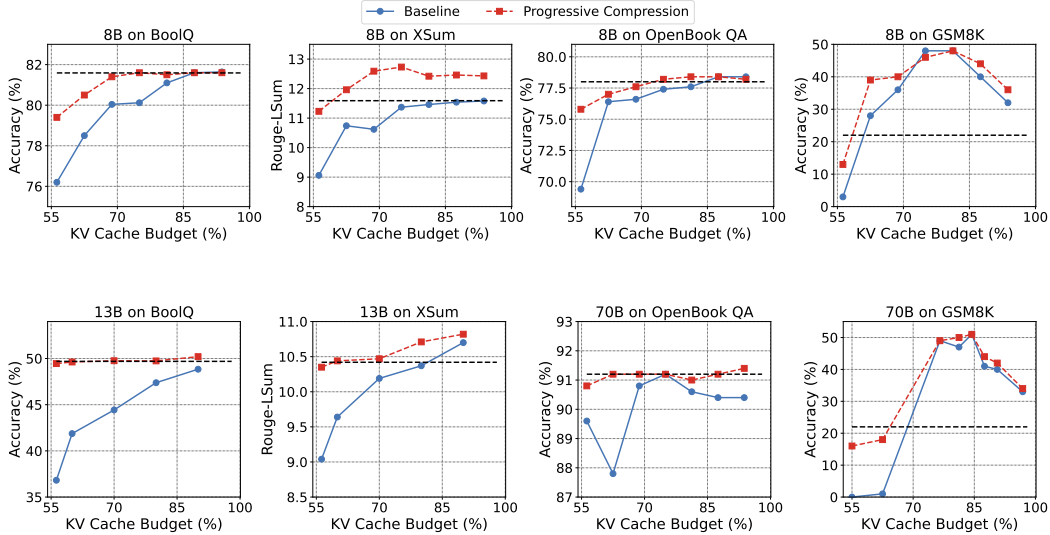



Figure 2: Performance of KV cache compression on LLaMA models. LORC compresses the KV weights with a progressive strategy, while the baselines compress each layer with the same ratio. The horizontal dashed line indicates the performance with a full-cache model.

313 merits make our method easily integrable in resource-constrained environments, enabling efficient
 314 model deployment with limited KV cache budgets.

315 In Figure 2, one interesting observation is that in some cases the model with a compressed KV cache
 316 leads to better performance. Particularly, on the GSM8K dataset, performing KV cache compression
 317 leads to more than 10% performance improvement. This phenomenon aligns with findings reported
 318 in the literature (Ge et al., 2023). Also, similar effects have been documented in the context of
 319 improving reasoning by applying low-rank decomposition on the MLP layers (Sharma et al., 2023).
 320 We believe this phenomenon demonstrates the feasibility of conducting task-specific profiling for
 321 better performance, or adapting our proposed method in model finetuning.

322 6.5 Single Layer Profiling



331 Figure 3: Single-layer compression results. This experiment
 332 uses LLaMA-3-Instruct-8B on the OpenBookQA dataset.
 333

To investigate the impact of compression at different layers, we conduct experiments on single-layer compression as shown in Fig. 3. We use LLaMA-3-Instruct-8B on OpenBook QA for this experiment. The original dimension of the KV head is 1024, and we select compression dimensions from [256, 384, 512] to compress each single layer while keeping all other layers untouched.

334 Figure 3 shows clear layer-specific variability, indicating that some layers are more susceptible to
 335 compression than others, particularly in the shallow layers. It is observed that the deep layers (*i.e.*,
 336 layers 15–31 of the 32-layer LLaMA-3-Instruct 8B model), despite the reduction in dimensions,
 337 maintain performance closely approaching the full KV Cache baseline. This suggests that these layers
 338 can sustain robust performance even when subjected to significant parameter reduction. This finding
 339 supports our progressive compression strategy for optimizing model efficiency without significantly
 340 compromising the model’s effectiveness.

Table 1: Performance comparison between compression on shallow layers and deep layers on OpenBookQA. For our progressive compression strategy, we report the performance at the 60% overall compression ratio. For layer-0 compression and shallow blocks compression, we use a 50% layer compression ratio within the chosen strategy. Hence, the overall compression ratio is 98.44% for the layer-0 compression, and 93.75% for the shallow blocks compression.

Model	Baseline	Ours	Layer 0	Shallow Blocks (1/8)
LLaMA-2-13b	76.6	77.4 ($\uparrow 0.8$)	77.2 ($\uparrow 0.6$)	74.8 ($\downarrow 1.8$)
LLaMA-3-Instruct-8b	78.0	77.4 ($\downarrow 0.6$)	67.2 ($\downarrow 10.8$)	61.4 ($\downarrow 16.6$)
LLaMA-3-Instruct-70b	91.2	91.2 ($\uparrow 0.0$)	84.2 ($\downarrow 7.0$)	23.2 ($\downarrow 68.0$)

341 6.6 Curse of shallow layers

342 To validate the intuition of the progressive compression strategy that the noise caused by shallow
 343 compressed layers will be amplified more after propagation, we compare it to compressing the first
 344 layer and the shallow blocks (i.e., the first 1/8 layers in a model) on 3 LLaMA models.

345 Table 1 shows how the compressed shallow layers impact the model performance, taking the baseline
 346 full-cache model and our method as reference. The results indicate that compressing only the first
 347 layer can lead to a performance decline, with reductions ranging from minimal to moderate. For
 348 instance, the LLaMA-3-70B gives a 7.0% decrease, while the LLaMA-3-Instruct-8b shows a more
 349 substantial drop of 10.8%. When compressing the shallow blocks, the impact is more pronounced.
 350 The LLaMA-3-Instruct-8B suffers a 16.6% reduction. Notably, the LLaMA-3-Instruct-70b model
 351 shows a drastic 68.0% decline, highlighting a significant sensitivity to shallow layer compression.

352 These findings underscore the importance of careful layer selection in compression strategies and
 353 validate the effectiveness of our progressive compression method, as the choice of layer to compress
 354 can have a substantial impact on model performance, particularly in larger or more complex models.

355 6.7 Memory footprint reduction analysis

Table 2: Summary of Model Sizes, KV cache usage and performance drop. Experiments were conducted with a batch size of 64 and a sequence length of 2048 for all models.

Model	KV Cache				Average Performance Drop	
	Full	dim	dim_c	Ours		Compression Ratio
LLaMA-2-13B	50G	5120	2048	27.5G	55%	0.47%
LLaMA-3-8B	8G	1024	512	4.8G	60%	0.92%
LLaMA-3-70B	20G	1024	512	11G	55%	0.22%

356 We report the memory footprint reduction in Table 2. By controlling the performance drop averaged
 357 on the four tasks less than 1%, we can achieve a considerable compression ratio from 55%-60%.
 358 For the LLaMA-3 models in which the GQA has already been employed to save the KV cache, we
 359 further achieve a significant compression ratio. Note that we have excluded the GSM8k results for
 360 the performance drop calculation for a fair comparison.

361 7 Conclusions

362 In conclusion, we proposed LORC, a novel approach to KV cache compression that capitalizes
 363 on the inherent low-rank properties of weight matrices. Our method employs a progressive layer-
 364 wise compression strategy, implementing a post-hoc low-rank approximation to circumvent the
 365 complexities and limitations associated with token-level eviction strategies and model retraining.
 366 Moreover, we provide theoretical analysis, deriving error bounds for layer compression and error
 367 propagation in deep networks, supporting our design of progressive compression strategy. This
 368 theoretically grounded and universally applicable approach preserves model integrity and performance
 369 across diverse tasks, attention mechanisms, and model scales. Our comprehensive experimental
 370 results demonstrate that LORC significantly reduces GPU memory requirements while minimally
 371 impacting performance. This approach offers a robust and efficient solution of KV cache compression,
 372 without requiring attention pattern analysis or model tuning.

373 References

- 374 Achiam, J., Adler, S., Agarwal, S., Ahmad, L., Akkaya, I., Aleman, F. L., Almeida, D., Altenschmidt,
375 J., Altman, S., Anadkat, S., et al. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023.
- 376 Ainslie, J., Lee-Thorp, J., de Jong, M., Zemlyanskiy, Y., Lebrón, F., and Sanghai, S. Gqa: Train-
377 ing generalized multi-query transformer models from multi-head checkpoints. *arXiv preprint*
378 *arXiv:2305.13245*, 2023.
- 379 Chowdhery, A., Narang, S., Devlin, J., Bosma, M., Mishra, G., Roberts, A., Barham, P., Chung,
380 H. W., Sutton, C., Gehrmann, S., et al. Palm: Scaling language modeling with pathways. *Journal*
381 *of Machine Learning Research*, 24(240):1–113, 2023.
- 382 Clark, C., Lee, K., Chang, M.-W., Kwiatkowski, T., Collins, M., and Toutanova, K. Boolq: Exploring
383 the surprising difficulty of natural yes/no questions. *arXiv preprint arXiv:1905.10044*, 2019.
- 384 Cobbe, K., Kosaraju, V., Bavarian, M., Chen, M., Jun, H., Kaiser, L., Plappert, M., Tworek, J.,
385 Hilton, J., Nakano, R., et al. Training verifiers to solve math word problems. *arXiv preprint*
386 *arXiv:2110.14168*, 2021.
- 387 de Jong, M., Zemlyanskiy, Y., Ainslie, J., FitzGerald, N., Sanghai, S., Sha, F., and Cohen, W.
388 Fido: Fusion-in-decoder optimized for stronger performance and faster inference. *arXiv preprint*
389 *arXiv:2212.08153*, 2022.
- 390 Ge, S., Zhang, Y., Liu, L., Zhang, M., Han, J., and Gao, J. Model tells you what to discard: Adaptive
391 kv cache compression for llms. *arXiv preprint arXiv:2310.01801*, 2023.
- 392 Guan, Y., Li, Z., Leng, J., Lin, Z., and Guo, M. Transkimmer: Transformer learns to layer-wise skim.
393 *arXiv preprint arXiv:2205.07324*, 2022.
- 394 Holmes, C., Tanaka, M., Wyatt, M., Awan, A. A., Rasley, J., Rajbhandari, S., Aminabadi, R. Y., Qin,
395 H., Bakhtiari, A., Kurilenko, L., et al. Deepspeed-fastgen: High-throughput text generation for
396 llms via mii and deepspeed-inference. *arXiv preprint arXiv:2401.08671*, 2024.
- 397 Izacard, G. and Grave, E. Leveraging passage retrieval with generative models for open domain
398 question answering. *arXiv preprint arXiv:2007.01282*, 2020.
- 399 Lin, C.-Y. ROUGE: A package for automatic evaluation of summaries. In *Text Summarization*
400 *Branches Out*, pp. 74–81, Barcelona, Spain, July 2004. Association for Computational Linguistics.
401 URL <https://www.aclweb.org/anthology/W04-1013>.
- 402 Liu, A., Feng, B., Wang, B., Wang, B., Liu, B., Zhao, C., Dengr, C., Ruan, C., Dai, D., Guo, D.,
403 et al. Deepseek-v2: A strong, economical, and efficient mixture-of-experts language model. *arXiv*
404 *preprint arXiv:2405.04434*, 2024a.
- 405 Liu, Z., Desai, A., Liao, F., Wang, W., Xie, V., Xu, Z., Kyrillidis, A., and Shrivastava, A. Scissorhands:
406 Exploiting the persistence of importance hypothesis for llm kv cache compression at test time.
407 *Advances in Neural Information Processing Systems*, 36, 2024b.
- 408 Liu, Z., Yuan, J., Jin, H., Zhong, S., Xu, Z., Braverman, V., Chen, B., and Hu, X. Kivi: A tuning-free
409 asymmetric 2bit quantization for kv cache. *arXiv preprint arXiv:2402.02750*, 2024c.
- 410 Mihaylov, T., Clark, P., Khot, T., and Sabharwal, A. Can a suit of armor conduct electricity? a new
411 dataset for open book question answering. *arXiv preprint arXiv:1809.02789*, 2018.
- 412 Mu, J., Li, X., and Goodman, N. Learning to compress prompts with gist tokens. *Advances in Neural*
413 *Information Processing Systems*, 36, 2024.
- 414 Narayan, S., Cohen, S. B., and Lapata, M. Don't give me the details, just the summary! topic-aware
415 convolutional neural networks for extreme summarization. *ArXiv*, abs/1808.08745, 2018.
- 416 Pope, R., Douglas, S., Chowdhery, A., Devlin, J., Bradbury, J., Heek, J., Xiao, K., Agrawal, S., and
417 Dean, J. Efficiently scaling transformer inference. *Proceedings of Machine Learning and Systems*,
418 5, 2023.

- 419 Ribar, L., Chelombiev, I., Hudlass-Galley, L., Blake, C., Luschi, C., and Orr, D. Sparq attention:
420 Bandwidth-efficient llm inference. *arXiv preprint arXiv:2312.04985*, 2023.
- 421 Sharma, P., Ash, J. T., and Misra, D. The truth is in there: Improving reasoning in language models
422 with layer-selective rank reduction. *ArXiv*, abs/2312.13558, 2023.
- 423 Shazeer, N. Fast transformer decoding: One write-head is all you need. *arXiv preprint*
424 *arXiv:1911.02150*, 2019.
- 425 Sheng, Y., Zheng, L., Yuan, B., Li, Z., Ryabinin, M., Chen, B., Liang, P., Ré, C., Stoica, I., and
426 Zhang, C. Flexgen: High-throughput generative inference of large language models with a single
427 gpu. In *International Conference on Machine Learning*, pp. 31094–31116. PMLR, 2023.
- 428 Su, J., Ahmed, M., Lu, Y., Pan, S., Bo, W., and Liu, Y. Roformer: Enhanced transformer with rotary
429 position embedding. *Neurocomputing*, 568:127063, 2024.
- 430 Sun, T., Liu, X., Zhu, W., Geng, Z., Wu, L., He, Y., Ni, Y., Xie, G., Huang, X., and Qiu, X. A simple
431 hash-based early exiting approach for language understanding and generation. *arXiv preprint*
432 *arXiv:2203.01670*, 2022.
- 433 Touvron, H., Martin, L., Stone, K., Albert, P., Almahairi, A., Babaei, Y., Bashlykov, N., Batra, S.,
434 Bhargava, P., Bhosale, S., et al. Llama 2: Open foundation and fine-tuned chat models. *arXiv*
435 *preprint arXiv:2307.09288*, 2023.
- 436 Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., and
437 Polosukhin, I. Attention is all you need. *Advances in neural information processing systems*, 30,
438 2017.
- 439 Yu, H., Yang, Z., Li, S., Li, Y., and Wu, J. Effectively compress kv heads for llm. *arXiv preprint*
440 *arXiv:2406.07056*, 2024.
- 441 Zhang, Y., Gao, B., Liu, T., Lu, K., Xiong, W., Dong, Y., Chang, B., Hu, J., Xiao, W., et al.
442 Pyramidkv: Dynamic kv cache compression based on pyramidal information funneling. *arXiv*
443 *preprint arXiv:2406.02069*, 2024a.
- 444 Zhang, Z., Sheng, Y., Zhou, T., Chen, T., Zheng, L., Cai, R., Song, Z., Tian, Y., Ré, C., Barrett,
445 C., et al. H2o: Heavy-hitter oracle for efficient generative inference of large language models.
446 *Advances in Neural Information Processing Systems*, 36, 2024b.
- 447 Zhou, W., Xu, C., Ge, T., McAuley, J., Xu, K., and Wei, F. Bert loses patience: Fast and robust
448 inference with early exit. *Advances in Neural Information Processing Systems*, 33:18330–18341,
449 2020.

450 **A Detailed Proofs**

451 **A.1 Proof of Theorem 1**

452 *The proof of Theorem 1 is presented here for completeness.*

453 **Proof.**

454 Let $W = U\Sigma V^\top$ be the full SVD of W , where $U \in \mathbb{R}^{m \times m}$ and $V \in \mathbb{R}^{n \times n}$ are orthogonal matrices,
455 and $\Sigma = \text{diag}(\sigma_1, \dots, \sigma_n)$ with singular values $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_n \geq 0$.

456 The rank- k approximation \tilde{W} is given by:

$$\tilde{W} = U_k \Sigma_k V_k^\top,$$

457 where U_k , Σ_k , and V_k are truncated versions of U , Σ , and V , respectively, keeping only the first k
458 singular values and corresponding vectors.

459 We have:

$$\begin{aligned} \|Wx - \tilde{W}x\|_2 &= \|(W - \tilde{W})x\|_2 \\ &= \|U(\Sigma - \Sigma_k)V^\top x\|_2 \\ &= \|(\Sigma - \Sigma_k)V^\top x\|_2, \quad \text{since } U \text{ is orthogonal} \\ &= \|\text{diag}(0, \dots, 0, \sigma_{k+1}, \dots, \sigma_n)V^\top x\|_2 \\ &\leq \sigma_{k+1}\|V^\top x\|_2 \\ &= \sigma_{k+1}\|x\|_2, \quad \text{since } V \text{ is orthogonal.} \end{aligned}$$

460

□

461 **A.2 Proof of Theorem 3**

462 *We present the proof of Theorem 3, including an adjustment for activation functions with Lipschitz*
463 *constant L_ϕ .*

464 **Proof.**

465 Let x_i and \tilde{x}_i denote the outputs of the i -th layer in the original and compressed networks, respectively.
466 We prove by induction that:

$$\|x_i - \tilde{x}_i\|_2 \leq \sum_{s=1}^i \left(\sigma_{k_s+1}^{(s)} L_\phi^{i-s} \prod_{j=s+1}^i \|W_j\|_2 \right). \quad (17)$$

467 **Base Case** ($i = 1$).

468 Using Theorem 1 and the Lipschitz property of ϕ :

$$\begin{aligned} \|x_1 - \tilde{x}_1\|_2 &= \|\phi(W_1 x_0) - \phi(\tilde{W}_1 x_0)\|_2 \\ &\leq L_\phi \|W_1 x_0 - \tilde{W}_1 x_0\|_2 \\ &\leq L_\phi \sigma_{k_1+1}^{(1)} \|x_0\|_2. \end{aligned}$$

469 **Inductive Step.**

470 Assume the inductive bound holds for layer $i - 1$. For layer i :

$$\begin{aligned}
\|x_i - \tilde{x}_i\|_2 &= \|\phi(W_i x_{i-1}) - \phi(\tilde{W}_i \tilde{x}_{i-1})\|_2 \\
&\leq L_\phi \|W_i x_{i-1} - \tilde{W}_i \tilde{x}_{i-1}\|_2 \\
&\leq L_\phi \left(\|W_i(x_{i-1} - \tilde{x}_{i-1})\|_2 + \|(W_i - \tilde{W}_i)\tilde{x}_{i-1}\|_2 \right) \\
&\leq L_\phi \left(\|W_i\|_2 \|x_{i-1} - \tilde{x}_{i-1}\|_2 + \sigma_{k_{i+1}}^{(i)} \|\tilde{x}_{i-1}\|_2 \right).
\end{aligned}$$

471 We can bound $\|\tilde{x}_{i-1}\|_2$ using the triangle inequality:

$$\|\tilde{x}_{i-1}\|_2 \leq \|x_{i-1}\|_2 + \|x_{i-1} - \tilde{x}_{i-1}\|_2.$$

472 Assuming that $\|x_{i-1}\|_2$ is bounded (which is reasonable in practice due to normalization techniques),
473 and applying the inductive hypothesis, we can express $\|x_i - \tilde{x}_i\|_2$ in terms of the accumulated errors
474 up to layer i .

475 By recursively applying this inequality and summing over all layers, we obtain the bound stated in
476 Theorem 3.

477 □

478 A.3 Note on Activation Functions and the Lipschitz Constant

479 It is important to note that Theorem 3 assumes the activation function ϕ has a Lipschitz constant L_ϕ ,
480 which reflects how much the function can amplify differences in its input. For activation functions like
481 ReLU, which are 1-Lipschitz, the error bound simplifies and indicates minimal error amplification
482 through the activation layers.

483 However, the LLaMA model family uses activation functions such as SwiGLU and GELU, whose
484 derivatives can exceed 1, making them not 1-Lipschitz. For networks employing such activation
485 functions, the error propagation bound in Theorem 2 is adjusted by incorporating a Lipschitz constant
486 L_ϕ , which may be greater than 1. This adjustment accounts for the potential additional error
487 amplification introduced by the activation functions.

488 B Reconstruction Error of Matrix SVD

489 Given a pre-trained LLM, we conduct layer-wise weight matrix decomposition and reconstruction.
490 We found that these matrices are low-rank and therefore can be reconstructed with low-dimension
491 matrices, resulting in minimal reconstruction error. It means instead complex eviction policy design
492 at the token level, we can turn to the attention level to develop a model and task agnostic KV cache
493 compression method. We present the relative reconstruction error in Figure 4, which is computed as
494 below.

```

495 1 # Code for matrix reconstruction error calculation
496 2 matrix_reconstructed = torch.mm(torch.mm(U_reduced, S_reduced_diag),
497   V_reduced)
498 3 error = torch.norm(matrix - matrix_reconstructed, p='fro')
499 4 relative_error = error / torch.norm(matrix, p='fro')
```

500 C Adjusted Position Embedding

501 Su et al. (2024) propose a rotary position embedding (RoPE) and it has been used in most recent
502 LLMs. Applying RoPE to self-attention gives

$$q_m^T k_n = (R_{\Theta, m}^d W_q^T x_m)^T (R_{\Theta, n}^d W_k^T x_n) = x^T W_q R_{\Theta, n-m}^d W_k^T x_n, \quad (18)$$

503 where Θ is a pre-defined rotary matrix, m and n denotes the token position. In practice, the rotation
504 matrix $R_{\Theta, n-m}^d$ is decomposed as $(R_{\Theta, m}^d)^T$ and $R_{\Theta, n}^d$ to rotate the query and key separately, and
505 the KV cache stores the rotated keys. To ensure that our compressed keys are compatible with the




Figure 4: Layerwise relative reconstruction errors. wk_{err} and wv_{err} denote the relative difference between the original key/value matrices and their corresponding low-rank approximations measured using the Frobinus norm. The compression ratio is computed as $r = \frac{d_c}{N_h \times d_h}$, where N_h is the number of attention heads and d_h, d_c is the original and compressed hidden dimensions respectively.

506 rotary operation, we adjust the position embedding pipeline. Specifically, we store the compressed
 507 keys $X(\Sigma V^T)_{D \times d_c}^T$ in cache, while incorporating the rotation and key projection into the query
 508 computation to streamline the process.

509 RoPE does bring additional computations to LoRC. To address this, we follow H2O (Zhang et al.,
 510 2024b) and FastGen (Ge et al., 2023) to develop a customized kernel that fuses cache reconstruction
 511 and rotation operations. This approach minimizes memory transfers and computational overhead,
 512 enhancing the overall throughput.

513 Specifically, we leverage Triton to implement a fused kernel that optimizes memory access and com-
 514 putation. With this kernel, we are able to combine key reconstruction from compressed representation
 515 and apply RoPE rotations on-the-fly. It also streamlines memory access because the single data load
 516 from global memory for compressed keys requires minimal data movement. In a complete RoPE
 517 workflow, there are two stages:

- 518 • **Caching:** We project keys to low-rank space without position embedding and store only the
 519 compressed representation.
- 520 • **Attention computation:** Within the customized kernel, we load compressed keys into shared
 521 memory, reconstruct and apply RoPE in a fused operation, and then compute attention scores
 522 with position-aware keys.

523 With these engineering efforts, we achieve higher throughput compared to baseline configurations.

Table 3: Model Architectures.

Model	Attention	Layers	Heads	KV Heads	Head Dimension	Model Dimension	Weight Shape
LLaMA-2-13B	MHA	40	40	40	128	5120	5120×5120
LLaMA-3-Instruct-8B	GQA	32	32	8	128	4096	4096×1024
LLaMA-3-Instruct-70B	GQA	80	64	8	128	8192	8192×1024

524 D Latency Analysis

Table 4: Latency analysis for sensitivity calculation and SVD processing.

Model	Computing Resource	Sensitivity Calculation (s)	SVD Processing (s)
LLaMA-3-8B-Instruct	NVIDIA H100 80GB HBM3 \times 1	30.6	33.3
LLaMA-3-70B-Instruct	NVIDIA H100 80GB HBM3 \times 8	12.2	14.3

525 The latency analysis is shown in Table 4. Note that sensitivity calculation only takes place once for a
 526 given model, and the SVD processing is a one-time implementation during the model initialization

527 stage. Such latency will not affect inference speed. Compared to the whole inference duration, the
 528 latency incurred by these processes is negligible. For example, the LLaMA-3-70B-Instruct model
 529 requires approximately 1 hour and 8 minutes to process 1,000 summaries from the XSUM dataset
 530 with a batch size of 32 and a sequence length of 8,000. The combined latency introduced by sensitivity
 531 calculations and SVD processing represents only 0.6% of the total inference time.

532 E Throughput Analysis

533 Our method involves additional computations to recover the compressed cache and manage RoPE.
 534 To address this, we developed a customized kernel that fuses cache reconstruction and rotation
 535 operations. This approach minimizes memory transfers and computational overhead, enhancing the
 536 overall throughput.

537 Following H2O () we conducted throughput experiments with fixed input and output sequence
 538 lengths using the LLaMA-3-70B-Instruct model on a node equipped with eight NVIDIA H100 80GB
 539 HBM3 GPUs. The results below indicate that our engineering efforts can streamline the attention
 540 computation with LoRC compression, thereby achieving higher throughput compared to full cache
 scenarios.

Table 5: Throughput of LLaMA-3-70B-Instruct with and without LoRC compression.

Input Length	Output Length	Batch Size	Full Cache Throughput	LoRC-60% Throughput	Speedup
1024	2048	32	52.75 tokens/s	60.08 tokens/s	$\times 1.14$
1024	4096	32	78.66 tokens/s	98.74 tokens/s	$\times 1.26$

541

542 F Comparison with Other KV Cache Compression Methods

543 We compare LoRC with a token-eviction method – H2O (Zhang et al., 2024b) and a quantization
 544 method – KIVI (Liu et al., 2024c) in this section. We conduct experiments using LLaMA-3-8B-
 545 Instruct. For H2O and LoRC, we keep the same 60% KV cache budget. For KIVI, we use the
 546 KIVI-4-bit implementation. We evaluate accuracy on BoolQ and OpenBook QA, and Rouge-Lsum
 547 for XSum. The results below show LoRC can preserve a better performance compared to token-
 548 eviction and quantization methods when KV cache is aggressively compressed.

Table 6: Performance comparison across different KV cache compression methods.

Method	BoolQ	XSum	OpenBookQA
Full Cache	81.6	11.6	78.0
H2O (Zhang et al., 2024b)	76.4	10.5	75.1
KIVI (Liu et al., 2024c)	77.6	10.3	74.8
LoRC	79.2	11.2	75.7

549 For throughput comparison with the other methods, we use the metric seconds per iteration (s/it)
 550 on XSum. We adhered to the same 60% KV cache budget for both H2O and LoRC. In an effort to
 551 demonstrate compatibility with other KV cache compression methods, we integrated LoRC with both
 552 H2O and standard 8-bit quantization. Specifically, we allocated a 70% KV cache budget to LoRC
 553 and 85% to H2O, resulting in approximately a 60% overall cache budget. A similar configuration
 was used for the combination of LoRC with 8-bit quantization.

Table 7: Throughput and performance comparison of different methods..

	Full Cache	H2O	KIVI	LoRC	LoRC w/ H2O	LoRC w/ 8-bit
Performance	11.6	10.5	10.3	11.2	11.0	10.9
s/it	228.5	192.7	186.6	203.4	195.9	193.7

554

555 Our results show that LoRC achieves superior performance compared to existing compression
 556 methods while maintaining competitive throughput. While the SVD-based computations introduce
 557 some overhead, LoRC still has a higher throughput than its full cache baseline because we deployed
 558 a customized kernel to streamline the attention computation. It is important to note that our primary
 559 objective was not to develop a Pareto-optimal method (simultaneously optimizing both performance
 560 and throughput), but rather to introduce an orthogonal approach to KV cache compression and
 561 establish the progressive compression strategy.

562 Additionally, the integration of LoRC with both token-eviction and quantization methods demonstrates
 563 its compatibility with existing compression methods. With only a slight performance trade-off
 564 (11.0/10.9 vs 11.2), these hybrid configurations achieve better throughput compared to using LoRC
 565 alone, which shows the potential of combining LoRC with orthogonal compression approaches.

566 G Enhancements in memory-constrained deployment scenarios

Table 8: LLaMA-3-8B-Instruct model on a single GPU with 24GB memory.

Max Sequence Length	Batch Size	8B w/ Full Cache	8B w/ LoRC-60%
2048	20	OOM	22.12 G

Table 9: LLaMA-3-70B-Instruct model on 8 GPUs with 24GB memory.

Max Sequence Length	Batch Size	70B w/ Full Cache	70B w/ LoRC-60%
2048	8	OOM	22.97 G

567 LoRC demonstrates significant benefits in memory-constrained deployment scenarios. To further
 568 validate this, we provide additional experiments here. These results demonstrate that LoRC enables
 569 the deployment of models that would otherwise be impossible to run with full KV cache on consumer
 570 GPUs.

571 H Prompts used in experiments

572 We present our prompts used for different datasets here. We use a few-shot setting for LLaMA-2
 573 models on OpenBookQA (1-shot), BoolQ (2-shot), and GSM8k (8-shot), and zero-shot setting for
 574 other experiments.

575 OpenBookQA

```

576 1 def format_examples(examples):
577 2     example_prompts = []
578 3     for j in range(1):
579 4         question = examples['question_stem'][j]
580 5         fact = examples['fact1'][j]
581 6         choices = examples['choices'][j]['text']
582 7         labels = examples['choices'][j]['label']
583 8         formatted_choices = "\n".join(f"{label} {text}" for label,
584         text in zip(labels, choices))
585 9         answer = examples['answerKey'][j]
586 0         example_prompt = f"Fact: {fact}\nQuestion: {question}\nOptions
587         :\n{formatted_choices}\nAnswer: {answer}\n"
588 1         example_prompts.append(example_prompt)
589 2     return "\n--\n".join(example_prompts)
590 3
591 4
592 5 def create_prompts_from_data(data, example_context):
593 6     prompts = []
594 7     answers = []
595 8     for i in range(len(data['id'])):
```

```

59619     question = data['question_stem'][i]
59720     fact = data['fact1'][i]
59821     choices = data['choices'][i]['text']
59922     labels = data['choices'][i]['label']
60023     formatted_choices = "\n".join(f"{label}) {text}" for label,
601     text in zip(labels, choices))
60224
60325     task_intro = "You will be provided with a fact and a related
604     question. Your task is to use the given fact to choose the correct
605     answer from the provided options."
60626     prompt = f"Task Introduction:\n{task_intro}\n1-Shot Examples:\n
607     n{example_context}\n---\nFact: {fact}\nQuestion: {question}\n
608     nOptions:\n{formatted_choices}\nAnswer:"
60927     prompts.append(prompt)
61028     answers.append(data['answerKey'][i])
61129     return prompts, answers
61230
61331
61432 def extract_option_label(outputs):
61533     answer_labels = []
61634     for output in outputs:
61735         match = re.search(r'\b([A-D])\b', output)
61836         if match:
61937             answer_labels.append(match.group(1))
62038         else:
62139             answer_labels.append(None)
62240     return answer_labels

```

623 BoolQ

```

624 1 def get_examples(dataset, num_examples):
625 2     selected_examples = dataset.shuffle(seed=42).select(range(
626     num_examples))
627 3     examples = []
628 4     for i in range(num_examples):
629 5         passage = selected_examples['passage'][i]
630 6         question = selected_examples['question'][i]
631 7         answer = "yes" if selected_examples['answer'][i] else "no"
632 8         examples.append((passage, question, answer))
633 9
63410     example_section = "\n\n".join([
63511         f"Example {i + 1}:\nPassage: {ex[0]}\nQuestion: {ex[1]}\n
636     nAnswer: {ex[2]}" for i, ex in enumerate(examples)
63712     ])
63813     return example_section
63914
64015 def create_prompts_from_data(data, example_section=None):
64116     task_description = "For each passage and question, determine if
642     the answer to the question is 'yes' or 'no' based on the passage
643     provided."
64417
64518     prompts = []
64619     references = []
64720     for question, passage, answer in zip(data['question'], data['
648     passage'], data['answer']):
64921         prompt = f"{task_description}\n\n2-Shot Examples:{
650     example_section}\n\nPassage: {passage}\nQuestion: {question}\n\
651     nAnswer (yes or no):"
65222         prompts.append(prompt)
65323         references.append("yes" if answer else "no")
65424     return prompts, references
65525
65626 def extract_answer(generated_text: str) -> str:
65727     normalized_text = generated_text.lower().strip()
65828     if normalized_text.startswith("yes"):

```

```

659:9         return "yes"
660:0     elif normalized_text.startswith("no"):
661:1         return "no"
662:2     return "unknown"

```

663 XSum

```

664 1 def create_prompts_from_data(data):
665 2     prompts = []
666 3     references = []
667 4     for article, summary in zip(data['document'], data['summary']):
668 5         prompt = f"Provide a concise summary of the text below: {
669         article}\n\nSummary:"
670 6         prompts.append(prompt)
671 7         references.append(summary)
672 8     return prompts, references

```

673 GSM8k

```

674 1 def create_prompts_from_data(data, examples):
675 2     content = f"Please give a step-by-step answer to the question. You
676         have to put your final numeric answer at the end, without any
677         extra sign, prefix, or suffix, just pure integer numbers, in the
678         format: \n#### answer\n Done, make sure to separate the final
679         numeric answer with \n####"
680 3
681 4     prompts = []
682 5     references = []
683 6
684 7     example_section = ""
685 8     for ex_question, ex_answer in examples:
686 9         example_section += f"\nExample Question: {ex_question}\n
687         nExample Answer: {ex_answer}\n"
688:0
689:1     for question, answer in zip(data['question'], data['answer']):
690:2         prompt = f"{example_section}\nQuestion: {question}\n{content}.
691         "
692:3         prompts.append(prompt)
693:4         _, extracted_answer = extract_answer(answer)
694:5         references.append(extracted_answer)
695:6     return prompts, references
696:7
697:8 def extract_answer(completion):
698:9     start_idx = completion.find("####")
699:0     if start_idx == -1:
700:1         return completion, 'None'
701:2     start_idx += 4 # Move past '####'
702:3     end_idx = completion.find('\n', start_idx)
703:4     if end_idx == -1:
704:5         end_idx = len(completion)
705:6     answer = completion[start_idx:end_idx].strip()
706:7     return completion[:end_idx], answer
707:8
708:9
709:0 def calculate_accuracy(predictions, references):
710:1     correct = sum([1 for (_, pred), ref in zip(predictions, references
711         ) if pred.lower() == ref.lower()])
712:2     return correct, len(predictions)

```

713 I Throughput Analysis

714 Our method involves additional computations to recover the compressed cache and manage RoPE.
715 To address this, we developed a customized kernel that fuses cache reconstruction and rotation

716 operations. This approach minimizes memory transfers and computational overhead, enhancing the
 717 overall throughput.

718 Following H2O () we conducted throughput experiments with fixed input and output sequence
 719 lengths using the LLaMA-3-70B-Instruct model on a node equipped with eight NVIDIA H100 80GB
 720 HBM3 GPUs. The results below indicate that our engineering efforts can streamline the attention
 721 computation with LoRC compression, thereby achieving higher throughput compared to full cache
 scenarios.

Table 10: Throughput of LLaMA-3-70B-Instruct with and without LoRC compression.

Input Length	Output Length	Batch Size	Full Cache Throughput	LoRC-60% Throughput	Speedup
1024	2048	32	52.75 tokens/s	60.08 tokens/s	× 1.14
1024	4096	32	78.66 tokens/s	98.74 tokens/s	× 1.26

722

723 J Comparison with Other KV Cache Compression Methods

724 We compare LoRC with a token-eviction method – H2O (Zhang et al., 2024b) and a quantization
 725 method – KIVI (Liu et al., 2024c) in this section. We conduct experiments using LLaMA-3-8B-
 726 Instruct. For H2O and LoRC, we keep the same 60% KV cache budget. For KIVI, we use the
 727 KIVI-4-bit implementation. We evaluate accuracy on BoolQ and OpenBook QA, and Rouge-Lsum
 728 for XSum. The results below show LoRC can preserve a better performance compared to token-
 729 eviction and quantization methods when KV cache is aggressively compressed.

Table 11: Performance comparison across different KV cache compression methods.

Method	BoolQ	XSum	OpenBookQA
Full Cache	81.6	11.6	78.0
H2O (Zhang et al., 2024b)	76.4	10.5	75.1
KIVI (Liu et al., 2024c)	77.6	10.3	74.8
LoRC	79.2	11.2	75.7

730 For throughput comparison with the other methods, we use the metric seconds per iteration (s/it)
 731 on XSum. We adhered to the same 60% KV cache budget for both H2O and LoRC. In an effort to
 732 demonstrate compatibility with other KV cache compression methods, we integrated LoRC with both
 733 H2O and standard 8-bit quantization. Specifically, we allocated a 70% KV cache budget to LoRC
 734 and 85% to H2O, resulting in approximately a 60% overall cache budget. A similar configuration
 was used for the combination of LoRC with 8-bit quantization.

Table 12: Throughput and performance comparison of different methods..

	Full Cache	H2O	KIVI	LoRC	LoRC w/ H2O	LoRC w/ 8-bit
Performance	11.6	10.5	10.3	11.2	11.0	10.9
s/it	228.5	192.7	186.6	203.4	195.9	193.7

735

736 Our results show that LoRC achieves superior performance compared to existing compression
 737 methods while maintaining competitive throughput. While the SVD-based computations introduce
 738 some overhead, LoRC still has a higher throughput than its full cache baseline because we deployed
 739 a customized kernel to streamline the attention computation. It is important to note that our primary
 740 objective was not to develop a Pareto-optimal method (simultaneously optimizing both performance
 741 and throughput), but rather to introduce an orthogonal approach to KV cache compression and
 742 establish the progressive compression strategy.

743 Additionally, the integration of LoRC with both token-eviction and quantization methods demonstrates
 744 its compatibility with existing compression methods. With only a slight performance trade-off
 745 (11.0/10.9 vs 11.2), these hybrid configurations achieve better throughput compared to using LoRC
 746 alone, which shows the potential of combining LoRC with orthogonal compression approaches.

747 **K Enhancements in memory-constrained deployment scenarios**

Table 13: LLaMA-3-8B-Instruct model on a single GPU with 24GB memory.

Max Sequence Length	Batch Size	8B w/ Full Cache	8B w/ LoRC-60%
2048	20	OOM	22.12 G

Table 14: LLaMA-3-70B-Instruct model on 8 GPUs with 24GB memory.

Max Sequence Length	Batch Size	70B w/ Full Cache	70B w/ LoRC-60%
2048	8	OOM	22.97 G

748 LoRC demonstrates significant benefits in memory-constrained deployment scenarios. To further
749 validate this, we provide additional experiments here. These results demonstrate that LoRC enables
750 the deployment of models that would otherwise be impossible to run with full KV cache on consumer
751 GPUs.