

SkipCat: Rank-Maximized Low-Rank Compression of Large Language Models via Shared Projection and Block Skipping

Yu-Chen Lu^{1,2}, Sheng-Feng Yu^{1,2}, Hui-Hsien Weng¹, Pei-Shuo Wang¹,
Yu-Fang Hu^{1,3}, Liang Hung-Chun³, Hung-Yueh Chiang⁴, Kai-Chiang Wu¹

¹National Yang Ming Chiao Tung University,

²Macronix International Co., Ltd.,

³Skymizer Taiwan Inc.,

⁴The University of Texas at Austin
yuchen.cs11@nycu.edu.tw

Abstract

Large language models (LLM) have achieved remarkable performance across a wide range of tasks. However, their substantial parameter sizes pose significant challenges for deployment on edge devices with limited computational and memory resources. Low-rank compression is a promising approach to address this issue, as it reduces both computational and memory costs, making LLM more suitable for resource-constrained environments. Nonetheless, naïve low-rank compression methods require a significant reduction in the retained rank to achieve meaningful memory and computation savings. For a low-rank model, the ranks need to be reduced by more than half to yield efficiency gains. Such aggressive truncation, however, typically results in substantial performance degradation. To address this trade-off, we propose *SkipCat*, a novel low-rank compression framework that enables the use of higher ranks while achieving the same compression rates. First, we introduce an intra-layer shared low-rank projection method, where multiple matrices that share the same input use a common projection. This reduces redundancy and improves compression efficiency. Second, we propose a block skipping technique that omits computations and memory transfers for selected sub-blocks within the low-rank decomposition. These two techniques jointly enable our compressed model to retain more effective ranks under the same compression budget. Experimental results show that, *without any additional fine-tuning*, our method outperforms previous low-rank compression approaches by 7% accuracy improvement on zero-shot tasks under the same compression rate. These results highlight the effectiveness of our rank-maximized compression strategy in preserving model performance under tight resource constraints.

1 Introduction

Large language models (LLM) (Yang et al. 2025; Touvron et al. 2023) have been employed in a wide range of real-world scenarios, such as professional problem-solving and intelligent home devices. However, deploying LLMs on edge devices remains challenging. For instance, the large number of parameters may exceed the limited memory capacity of such devices. In addition, their high computational cost can lead to excessive latency and energy consumption,

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

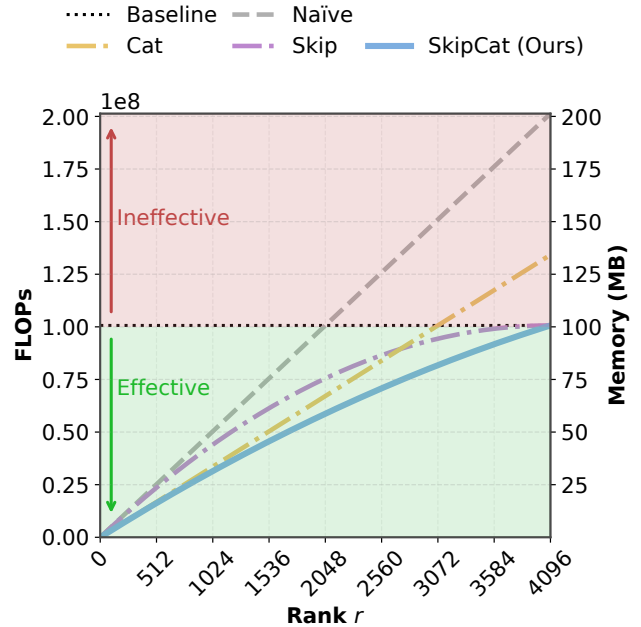


Figure 1: Computational cost (FLOPs) and memory footprint (MB) of the LLaMA2-7B attention module versus retained rank r . The black dashed line represents the original computational cost and memory footprint of the module. The gray dash-dot line corresponds to a naïve SVD compression, which only achieves effective compression when the retained rank is reduced to less than half of the full rank. In contrast, *SkipCat* maximizes the ranks and reduces the computational and memory cost, resulting in better trade-offs between efficiency and performance.

which limits their practical use in resource-constrained environments. Consequently, substantial research (Yuan et al. 2024) efforts have been devoted to enhancing the efficiency of LLM inference.

Model compression (Zhu et al. 2024) has become a widely used strategy to reduce both model size and computational cost. Quantization techniques achieve this by lowering the precision of both model parameters and activations

(Frantar et al. 2022; Liu et al. 2024), thereby reducing memory transfer overhead and potentially improving computational efficiency. Model pruning (Akhauri et al. 2024; Sun et al. 2023) improves inference efficiency by removing redundant or less important parameters.

In addition to the aforementioned techniques, low-rank compression (Yuan et al. 2023; Wang et al. 2025b) has also emerged as a promising research direction. By factorizing a weight matrix into two matrices with a shared low rank, we can reduce the number of parameters. Lowering the rank increases the compression rate by reducing the number of parameters and operations, but this may come at the cost of reduced accuracy. Conversely, a higher rank allows the model to better preserve its original performance, at the cost of a larger model size. A well-known limitation of existing low-rank compression methods often requires a substantial decrease in ranks to obtain efficiency gains, which can significantly impair model performance. As illustrated in Figure 1, taking the attention module of LLaMA2-7B as an example, the existing methods only begin to reduce floating-point operations (FLOPs) and memory below that of the original module when the retained rank falls below half of the original value (i.e., full rank). An excessive reduction can result in a significant drop in accuracy. Motivated by this observation, we aim to develop a method that increases the number of retained ranks while achieve the same compression rates, thereby mitigating the loss in model performance. To this end, we propose *SkipCat*, a low-rank compression method that maximizes the preserved ranks than existing approaches under the same compression rate, thereby allowing the model to retain better performance after compression.

Unlike naïve SVD methods that decompose each weight matrix separately, we propose *Intra-Layer Shared Low-Rank Projection using Matrix Concatenation (Cat)*, as illustrated in Figure 2(c). This technique constitutes one of the core components of our framework. In *Cat*, multiple matrices with the same input share a single projection matrix, reducing the number of projection components and enabling the model to retain a greater number of ranks under the same compression budget. Another key technique in our approach is *Block Skipping (Skip)*. As illustrated in Figure 2(d), *Skip* enables the omission of sub-block computations across all low-rank projections, allowing additional ranks to be preserved without increasing the compression budget. Our method, *SkipCat*, substantially improves rank retention and achieves highly effective compression rates, as shown in the red curve in Figure 1.

We evaluate our method on both the LLaMA (Touvron et al. 2023) and Qwen (Yang et al. 2025) models. At an equivalent compression rate, *SkipCat* outperforms existing low-rank compression methods in the *training-free* setting, achieving lower perplexity and up to a 7% improvement in zero-shot task accuracy. These results demonstrate that both of our techniques (i.e., *Skip* and *Cat*) enable the maximization of effective ranks while maintaining efficiency gains and delivering superior model performance at higher compression levels. We further evaluate *SkipCat* on larger-scale models and in conjunction with LoRA fine-tuning techniques. The consistently strong results underscore the gener-

alizability of our approach. Our approach offers a practical and effective solution for compressing LLMs, making them more feasible for edge devices. Our findings open new possibilities for future work in efficient model compression and deployment of large models in constrained environments. Our main contributions are summarized as follows:

- We propose a novel low-rank compression framework aimed at maximizing the number of effective ranks. The framework introduces two key techniques: (1) *Cat*, which enables multiple matrices to share a common low-rank projection; and (2) *Skip*, which allows the model to skip the computation of certain submatrices, while also addressing numerical instability issues that arise in half-precision inference.
- Experimental results show that our method consistently outperforms existing low-rank compression approaches in terms of compression performance.

2 Related Works

Numerous studies have explored techniques to reduce the computational and memory overhead of large language model inference. One major line of research focuses on lowering numerical precision, known as quantization. GPTQ (Frantar et al. 2022) is a state-of-the-art weight-only quantization method that converts model weights to low-precision formats, thereby reducing memory transfer costs during inference. Other works address weight-activation quantization, where the main challenge lies in handling activation outliers, which can severely degrade quantization accuracy. To address this, methods such as SmoothQuant (Xiao et al. 2023) and SpinQuant (Liu et al. 2024) apply different scaling strategies to mitigate the impact of outliers, effectively preserving model performance after quantization. When supported by hardware, activation quantization can also lead to faster inference by enabling low-precision arithmetic throughout the computation.

In addition to reducing bit-widths, another major line of research focuses on minimizing the number of parameters to improve both memory and computation efficiency. Pruning techniques (Sun et al. 2023; Frantar and Alistarh 2023) aim to remove redundant and less important parameters from the model. While unstructured pruning can achieve high sparsity, it typically requires specialized hardware to realize performance gains. As a result, structured pruning (Ma, Fang, and Wang 2023), which removes entire units such as attention heads or channels, is more hardware-friendly.

Low-rank compression is another promising approach for reducing the number of model parameters. A common approach employs singular value decomposition (SVD) to approximate weight matrices in a lower-dimensional space. In the context of LLMs, several works have explored SVD-based techniques. LORD (Kaushal, Vaidhya, and Rish 2023) applies SVD to compress code-focused LLMs, though directly applying naïve SVD to LLMs can lead to significant performance degradation. To mitigate such losses, FWSVD (Hsu et al. 2022) leverages gradient information to assess each parameter’s importance, thereby minimizing the impact of compression on model accuracy. On the other hand,

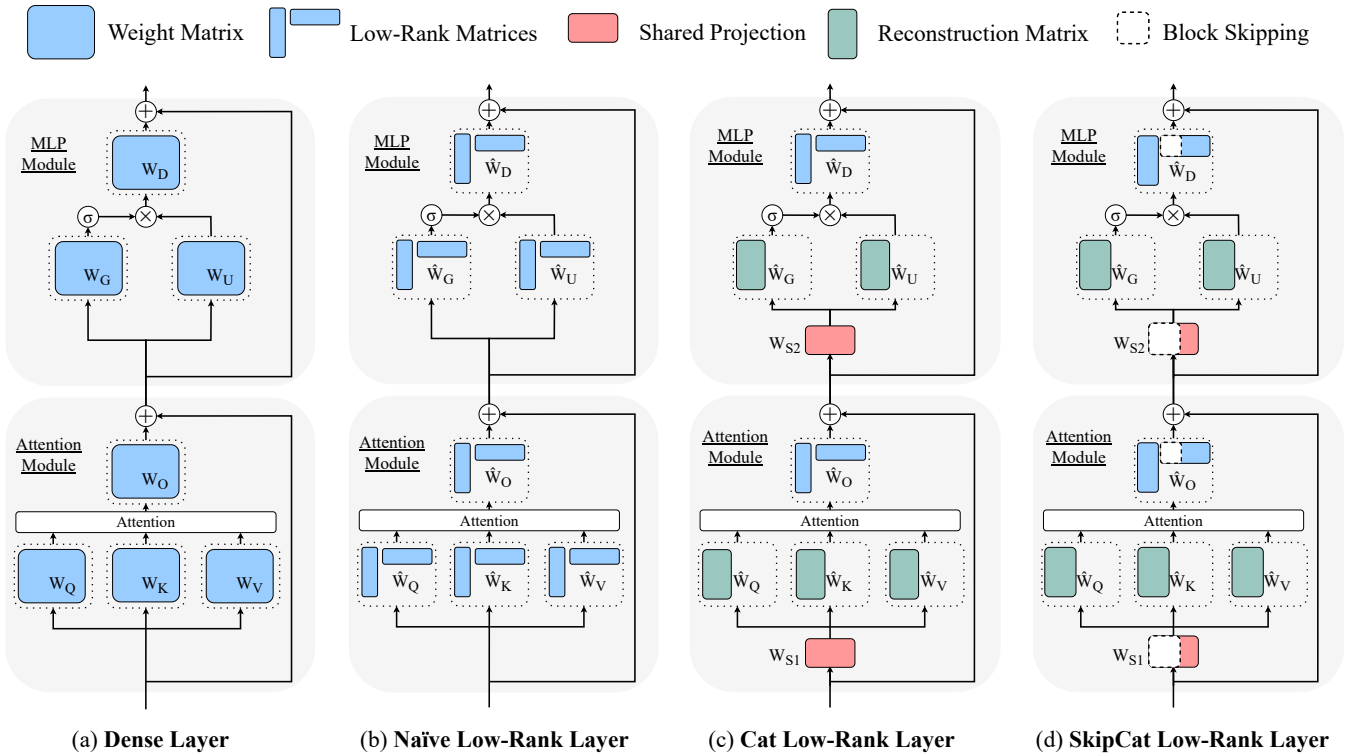


Figure 2: An illustration comparing our proposed method *SkipCat* with naïve low-rank compression. (a) The original model layer. (b) The model structure after applying naïve SVD low-rank compression. (c) Our *Cat* technique, which shares low-rank projection matrices across modules, enabling more ranks to be preserved without changing the compression rate. (d) Our *Skip* technique, applied on top of the *Cat* architecture, skips submatrix computations for all low-rank projections and further increases the number of preserved ranks.

the presence of outliers in LLM activations not only hinders the effectiveness of quantization, but also poses a significant challenge for low-rank compression. ASVD (Yuan et al. 2023) identifies this issue and proposes to mitigate the impact of outliers by applying a scaling matrix, derived from the activation distribution on calibration data. LatentLLM (Koike-Akino et al. 2025) enhances compression performance through the use of junction matrices. However, directly applying these matrices may lead to numerical instability. SVD-LLM (Wang et al. 2025b,a) applies a whitening transformation to the parameter matrices as a preprocessing step, making the performance degradation approximately proportional to the discarded singular values. This leads to better performance compared to directly applying SVD without whitening. In contrast to prior methods, our proposed approach aims to increase the number of retained ranks without altering the overall compression rate, and simultaneously resolves the numerical instability that arises during low-precision computation.

Low-rank compression is also compatible with other techniques and can be integrated to achieve greater efficiency gains. Dobi-SVD (Qinsi et al. 2025) integrates quantization with low-rank compression to achieve higher compression ratios without significant loss in model performance. Basis Sharing (Wang et al. 2024) shares the basis matrix among

low-rank modules of the same projection type in different layers. While the above-mentioned method helps reduce the model’s storage footprint in main memory, it brings no benefit in terms of computational cost or memory transfer during inference. In contrast, our method shares the projection matrix within *the same layer* for matrices that operate on the same input. This design not only reduces memory transfer overhead but also alleviates computational cost.

3 Proposed Method

In this section, we first introduce the background of SVD-based compression in Section 3.1. We then present the two core components of our method in Sections 3.2 and 3.3. Finally, Section 3.4 integrates these components into a unified framework and discusses the resulting benefits.

3.1 SVD-based Low-Rank Compression

The typical layer structure of LLM, comprising attention and MLP modules, is illustrated in Figure 2(a). For LLM, reducing the FLOPs of matrix computations and the size of weight matrices can help decrease computation time (Kaplan et al. 2020) or memory transfer latency during both the prefilling and decoding stages. Given a weight matrix $W \in \mathbb{R}^{d_{out} \times d_{in}}$, and an input vector $x \in \mathbb{R}^{d_{in} \times 1}$, the output is $y = Wx \in \mathbb{R}^{d_{out} \times 1}$. We assume a batch size of

1 for simplicity. Each of the d_{out} elements in y requires d_{in} multiplications and d_{in} additions, leading to a total of $2d_{\text{in}}d_{\text{out}}$ FLOPs. Additionally, storing the weight matrix requires $d_{\text{in}}d_{\text{out}}$. This setting corresponds to the baseline case, as indicated by the black dashed line in Figure 1.

Low-rank compression approximates large weight matrices with the product of smaller matrices, leading to reductions in both memory consumption and computation. The naïve approach is typically grounded in singular value decomposition (SVD), which factorizes a matrix into a set of orthogonal components. The weight matrix $W \in \mathbb{R}^{d_{\text{out}} \times d_{\text{in}}}$ can be decomposed using SVD as follows:

$$W = U\Sigma V^\top,$$

where $U \in \mathbb{R}^{d_{\text{out}} \times d_{\text{out}}}$ and $V \in \mathbb{R}^{d_{\text{in}} \times d_{\text{in}}}$ are orthogonal matrices. The matrix $\Sigma \in \mathbb{R}^{d_{\text{out}} \times d_{\text{in}}}$ is a rectangular diagonal matrix, i.e. the top-left $R \times R$ block is a diagonal matrix, where $R = \text{rank}(W) \leq \min(d_{\text{in}}, d_{\text{out}})$, and its nonzero diagonal entries $\sigma_1, \sigma_2, \dots, \sigma_R$ are the singular values of W , arranged in non-increasing order:

$$\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_R > 0.$$

Each singular value indicates how much its associated component contributes to reconstructing the original matrix. Therefore, to achieve the best approximation at a given compression rate, one can retain only the top- r singular values with $r \leq R$ and their corresponding components from the decomposition. This yields a truncated diagonal matrix $\Sigma_r \in \mathbb{R}^{r \times r}$ containing the leading singular values. Simultaneously, we retain only the first r columns of U and V , resulting in the truncated matrices $U_r \in \mathbb{R}^{d_{\text{out}} \times r}$ and $V_r \in \mathbb{R}^{d_{\text{in}} \times r}$, respectively. The matrix W can then be approximated by a rank- r reconstruction:

$$W \approx U_r \Sigma_r V_r^\top.$$

To reduce computational cost during LLM inference, the weight-input multiplication can be approximated using the rank- r decomposition. By absorbing the singular values into the orthogonal matrices, the weight matrix can be expressed as the product of two smaller matrices, without increasing computational overhead, that is

$$Wx \approx BAx, \quad (1)$$

where $B = U_r \Sigma_r^{1/2} \in \mathbb{R}^{d_{\text{out}} \times r}$ and $A = \Sigma_r^{1/2} V_r^\top \in \mathbb{R}^{r \times d_{\text{in}}}$. We define A as the projection matrix and B as the reconstruction matrix in the low-rank decomposition, where A projects the input into a lower-dimensional space and B reconstructs the output. Existing low-rank compression methods (Wang et al. 2025b, 2024) typically decompose the major weight matrices within each layer of the LLM, as illustrated in Figure 2(b).

After applying low-rank compression with rank r , the total number of parameters in matrices A and B is $r(d_{\text{in}} + d_{\text{out}})$, with a computational cost of $2r(d_{\text{in}} + d_{\text{out}})$ FLOPs. Consequently, low-rank compression leads to a computational gain only when the reduced FLOPs and parameters are fewer than the original matrix multiplication. This condition yields the constraint:

$$r < d_{\text{in}}d_{\text{out}} / (d_{\text{in}} + d_{\text{out}}). \quad (2)$$

Equation (2) highlights the limitation under which low-rank compression begins to yield computational benefits. For example, if the weight matrix W is a full-rank square matrix, i.e. $R = d_{\text{in}} = d_{\text{out}}$. The naïve low-rank compression becomes efficient only when the rank r is less than half the full rank R , that is $r < R/2$, as illustrated by the gray dash-dot line in Figure 1. This strong constraint motivates our work, which explores alternative approaches to relax this limitation and enable more efficient compression.

3.2 Cat: Intra-Layer Shared Low-Rank Projection using Matrix Concatenation

As shown in the model architecture in Figure 2(b), we observe that in the attention module, q , k , and v matrices share the same input, while in the MLP module, $gate$ and up matrices also receive the same input. Therefore, if we allow matrices with identical inputs to share a single low-rank projection, as illustrated by the red components in Figure 2(c), we can increase the number of retained ranks without changing the overall compression ratio. To enable multiple modules to share the same low-rank projection, we apply *concatenated* SVD, thereby obtaining a shared low-rank projection. We illustrate the following examples using the three matrices in the attention module, namely W_Q , W_K , and W_V , each in $\mathbb{R}^{d_{\text{out}} \times d_{\text{in}}}$. We first concatenate these matrices along the output dimension, as follows:

$$W_{QKV} = [W_Q^\top \quad W_K^\top \quad W_V^\top]^\top \in \mathbb{R}^{3d_{\text{out}} \times d_{\text{in}}},$$

We then apply low-rank factorization to the concatenated matrix W_{QKV} decomposing it into two low-rank matrices B_{QKV} and A_{QKV} as follows:

$$W_{QKV}x \approx B_{QKV}A_{QKV}x = [\hat{W}_Q^\top \quad \hat{W}_K^\top \quad \hat{W}_V^\top]^\top W_{S1}x,$$

where $B_{QKV} \in \mathbb{R}^{3d_{\text{out}} \times r}$ and $A_{QKV} \in \mathbb{R}^{r \times d_{\text{in}}}$. We then slice B_{QKV} along the output dimension to recover the original shapes corresponding to \hat{W}_Q , \hat{W}_K , and \hat{W}_V which serve as the reconstruction matrices in the attention module, as shown in Figure 2(c). The A_{QKV} corresponds to the shared projection $W_{S1} \in \mathbb{R}^{r \times d_{\text{in}}}$. This approach can also be applied to *gate* and *up* matrices in MLP modules.

With low-rank compression via shared projection, the amortized number of parameters per matrix is reduced to $r(d_{\text{in}} + Cd_{\text{out}})/C$, and the corresponding FLOPs per matrix become $2r(d_{\text{in}} + Cd_{\text{out}})/C$, where C denotes the number of concatenated matrices. We refer to the above method as *Cat*, as illustrated by the orange line in Figure 1. Compared to the naïve approach, *Cat* allows more ranks to be enabled the retention of more ranks under the same memory and computational budget.

3.3 Skip: Block Skipping via Schur Complement and Column Permutation

In this subsection, we introduce a novel technique termed Block Skipping (*Skip*), inspired by the Schur complement. Specifically, the multiplication of the weight matrix and the input vector is approximated using two low-rank matrices,

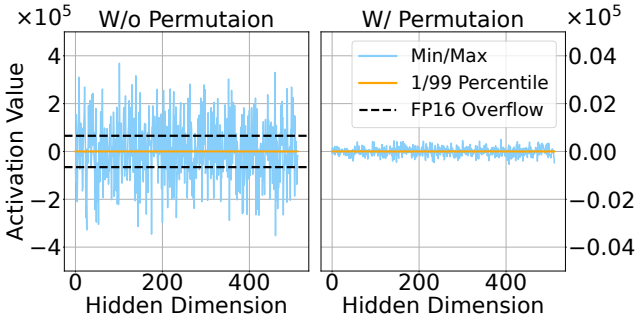


Figure 3: The output activation distribution of low-rank projections in LLaMA2-7B after applying block skipping. Left: Without column permutation, activation outliers cause FP16 overflow. Right: Column permutation stabilizes the distribution, preventing overflow in low-precision inference.

a low-rank projection matrix and a reconstruction matrix, as defined in Equation (1). We restructure the low-rank projection matrix into two subblocks and merge the first block into the reconstruction matrix, reducing the overall computational complexity. The core idea behind block skipping is to bypass a designated submatrix during computation, significantly reducing cost while maintaining the stability of the computation.

In detail, by analyzing the multiplication of the low-rank projection matrix $A \in \mathbb{R}^{r \times d_{in}}$ with the input vector $x \in \mathbb{R}^{d_{in} \times 1}$. We partition the matrix and vector as:

$$A = [A_1 \quad A_2], \quad x^\top = [x_1^\top \quad x_2^\top]^\top$$

where $A_1 \in \mathbb{R}^{r \times r}$, $A_2 \in \mathbb{R}^{r \times (d_{in}-r)}$, $x_1 \in \mathbb{R}^{r \times 1}$, and $x_2 \in \mathbb{R}^{(d_{in}-r) \times 1}$. Then, by assuming A_1 is invertible, the low-rank approximation can be rewritten as:

$$\begin{aligned} Wx \approx BAx &= B [A_1 \quad A_2] \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} \\ &= B(A_1x_1 + A_2x_2) \\ &= BA_1(x_1 + A_1^{-1}A_2x_2) \\ &= B'(x_1 + A'x_2) \end{aligned} \quad (3)$$

by defining $B' = BA_1 \in \mathbb{R}^{d_{out} \times r}$ and $A' = A_1^{-1}A_2 \in \mathbb{R}^{r \times (d_{in}-r)}$.

This reformulation eliminates the need to explicitly compute products involving A_1 by absorbing it into the redefined matrices B' and A' . As a result, computing the expression in Equation (3) requires $2r(d_{in} + d_{out} - r) + r$ FLOPs, accounting for the multiplication with A' , vector addition, and the final projection via B' . The number of parameters is reduced to $r(d_{in} + d_{out} - r)$. As shown in Figure 1, *Skip* also retains more ranks than the naïve low-rank compression under the same compression rate.

Despite the promising efficiency of Equation (3), the formulation may suffer from *numerical instability* when A_1 is ill-conditioned. In such cases, A_1^{-1} can have large values, which in turn amplify the magnitude of A' . This amplification can lead to overflow issues when computing $A'x_2$ under

FP16 precision. As shown in the left subfigure of Figure 3, although most activation values remain small, some outliers exceed the maximum representable value of FP16, resulting in overflow. Consequently, this prevents inference from being performed reliably under low-precision settings.

To address this, we apply a column permutation P to the matrix A , such that the leading r columns of the permuted matrix $\tilde{A} = AP = [\tilde{A}_1 \quad \tilde{A}_2]$ form a *well-conditioned* submatrix $\tilde{A}_1 \in \mathbb{R}^{r \times r}$, with the remaining columns denoted by $\tilde{A}_2 \in \mathbb{R}^{r \times (d_{in}-r)}$. The Strong Rank-Revealing QR factorization (Gu and Eisenstat 1996) provides an effective algorithm to identify such well-conditioned column subsets, with provable bounds on condition number. Incorporating this permutation, the original formulation is rewritten as:

$$\begin{aligned} Wx \approx BAx &= BAPP^\top x = B\tilde{A}\tilde{x} \\ &= B [\tilde{A}_1 \quad \tilde{A}_2] \begin{bmatrix} \tilde{x}_1 \\ \tilde{x}_2 \end{bmatrix} \\ &= B\tilde{A}_1(\tilde{x}_1 + \tilde{A}_1^{-1}\tilde{A}_2\tilde{x}_2) \\ &= \tilde{B}'(\tilde{x}_1 + \tilde{A}'\tilde{x}_2) \end{aligned} \quad (4)$$

where $\tilde{x} = P^\top x$, $\tilde{B}' = B\tilde{A}_1$, and $\tilde{A}' = \tilde{A}_1^{-1}\tilde{A}_2$.

The permutation enhances numerical stability by avoiding inversion of poorly conditioned matrices and improves robustness in low-rank approximation. As shown in the right subfigure of Figure 3, our preprocessing step applies a column permutation to the weight matrix, which significantly stabilizes the activation values after computing $\tilde{A}'\tilde{x}_2$. Compared to the left subfigure, the resulting activations are nearly two orders of magnitude smaller and more uniformly distributed. This contribution enables block skipping to be stably executed under FP16 inference without incurring performance degradation.

3.4 SkipCat: Efficient Shared Projection with Block Skipping

In our work, our primary objective is to increase the number of retained ranks without changing the overall compression ratio. In the preceding subsections, we proposed two methods that individually enhance rank retention under the same budget. Building on these, we further integrate both techniques, as shown in Figure 2(d), where all low-rank projections are equipped with block skipping. The benefit of combining both methods can be observed in Figure 1. The green curve corresponds to applying the *Skip* technique to the low-rank modules. Compared to the standard approach (gray dash-dot line), our method substantially increases the number of retained ranks in the low-compression regime (i.e., within the $0.75 \sim 1 \times 10^8$ FLOPs or $75 \sim 100$ MB range). While the advantage of preserving more ranks diminishes under higher compression rates, we address this limitation by integrating the *Skip* and *Cat* techniques into a unified approach, referred to as *SkipCat*, as illustrated by the red line in Figure 1. Our proposed method successfully addresses the limitations of conventional low-rank compression. As illustrated by the blue line in Figure 1, *SkipCat* consistently operates within the effective compression region,

| Model | Comp. Rate | Method | Perplexity ↓ | | Zero-shot Task Accuracy (%) ↑ | | | | | | | Avg. (%) ↑ | Drop (%) ↓ |
|-----------|------------|---------------|--------------|--------------|-------------------------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|-------------|
| | | | Wiki2 | C4 | ARC-e | ARC-c | Hella | OBQA | Wino | MathQA | PIQA | | |
| LLaMA2-7B | - | Dense | 5.47 | 7.26 | 76.30 | 43.34 | 57.14 | 31.40 | 69.14 | 28.17 | 78.07 | 54.79 | - |
| | 20% | ASVD | 9.06 | 11.66 | 67.80 | 34.56 | 48.25 | 29.60 | 63.77 | 25.49 | 72.20 | 48.81 | 5.98 |
| | | Basis Sharing | 9.39 | 23.30 | 59.18 | 27.82 | 38.22 | 24.20 | 66.54 | 23.89 | 66.54 | 43.77 | 11.02 |
| | | Dobi-SVD | 9.39 | 19.46 | 54.25 | 23.29 | 38.80 | 22.80 | 58.56 | 22.75 | 65.34 | 40.83 | 13.97 |
| | | SVD-LLM | 8.82 | 13.42 | 58.67 | 27.65 | 43.10 | 26.20 | 64.25 | 23.85 | 70.18 | 44.84 | 9.95 |
| | | SkipCat | 6.29 | 8.95 | 73.23 | 40.02 | 51.90 | 31.00 | 68.43 | 27.71 | 75.84 | 52.59 | 2.20 |
| | 30% | ASVD | 208.55 | - | 34.89 | 23.12 | 29.89 | 13.40 | 51.62 | 22.08 | 56.86 | 33.12 | 21.67 |
| | | Basis Sharing | 12.47 | 38.81 | 51.22 | 22.95 | 34.25 | 19.60 | 59.43 | 23.82 | 61.92 | 39.03 | 15.77 |
| | | SVD-LLM | 11.75 | 19.37 | 52.90 | 24.06 | 38.01 | 22.00 | 61.72 | 22.88 | 67.19 | 41.25 | 13.54 |
| | | SkipCat | 7.65 | 11.57 | 68.31 | 35.67 | 45.95 | 26.20 | 65.98 | 24.66 | 72.42 | 48.46 | 6.34 |
| Dense | | 9.72 | 15.42 | 83.54 | 55.97 | 57.13 | 31.00 | 67.80 | 49.61 | 76.88 | 60.28 | - | |
| Qwen3-8B | 20% | ASVD | 22.10 | 35.42 | 73.99 | 41.72 | 42.02 | 26.40 | 63.30 | 31.42 | 72.09 | 50.13 | 10.14 |
| | | SVD-LLM | 14.33 | 23.21 | 71.38 | 44.54 | 47.29 | 28.80 | 66.46 | 30.99 | 72.14 | 51.66 | 8.62 |
| | | SkipCat | 11.68 | 19.09 | 78.96 | 49.49 | 52.47 | 29.60 | 69.38 | 40.37 | 74.65 | 56.42 | 3.86 |
| | 30% | ASVD | 108.80 | 138.07 | 47.98 | 23.89 | 30.57 | 17.60 | 50.99 | 23.08 | 61.64 | 36.54 | 23.74 |
| | | SVD-LLM | 19.17 | 32.69 | 58.08 | 33.28 | 40.18 | 23.60 | 61.72 | 25.33 | 68.23 | 44.35 | 15.93 |
| | | SkipCat | 13.81 | 23.70 | 74.28 | 42.41 | 46.57 | 28.80 | 65.43 | 32.16 | 72.31 | 51.71 | 8.57 |

Table 1: Perplexity and zero-shot accuracy of low-rank compression methods.

demonstrating that the retained ranks contribute proportionally to compression efficiency. This indicates that, under our approach, increasing the number of preserved ranks directly translates into meaningful compression gains. *SkipCat* effectively maximizes the number of effective ranks within a given computational and memory budget.

4 Experiments

4.1 Experimental Setup

Models and Datasets We evaluate *SkipCat* and previous low-rank compression methods on models from the LLaMA (Touvron et al. 2023) and Qwen (Yang et al. 2025) families. Since most previous compression methods have been evaluated on LLaMA2, we compare performance using both the 7B and 13B variants. To assess the generalizability of our approach beyond LLaMA-based models, we further evaluate it on Qwen3-8B and Qwen3-14B. We evaluate perplexity on the WikiText-2 (Merity et al. 2016) and C4 (Rafel et al. 2020) datasets. In addition, we assess zero-shot task accuracy using the LM-Evaluation-Harness framework (Gao et al. 2021) on a set of benchmark tasks, including ARC-Easy, ARC-Challenge (Clark et al. 2018), HellaSwag (Zellers et al. 2019), OpenBookQA (Mihaylov et al. 2018), WinoGrande (Sakaguchi et al. 2019), MathQA (Amini et al. 2019) and PIQA (Bisk et al. 2020).

Experimental Details We compare *SkipCat* against several well-known low-rank compression approaches, including ASVD (Yuan et al. 2023), Basis Sharing (Wang et al. 2024), Dobi-SVD (Qinsi et al. 2025), and SVD-LLM (Wang et al. 2025b). Following previous works, our reported compression rates are based on the memory transfer cost associated with the core weight matrices: q , k , v , o , $gate$, up , and $down$. We use the official codebases of ASVD and Basis Sharing¹, and follow their original settings to reproduce the

¹Compression rate for Basis Sharing is recalculated using our standard, differing from their original setup.

results. We evaluate Dobi-SVD using its officially released “noremapping” model checkpoints without further modification. For SVD-LLM and *SkipCat*, we perform whitening (Wang et al. 2025b) on the models using samples from the training splits of WikiText-2 and C4.

| Model | Comp. Rate | Method | Zero-shot Task Accuracy | |
|------------|------------|--------------|-------------------------|-------------|
| | | | Avg. (%) ↑ | Drop (%) ↓ |
| LLaMA2-13B | - | Dense | 58.07 | - |
| | 20% | SVD-LLM | 50.09 | 7.98 |
| | | SkipCat | 54.15 | 3.92 |
| 30% | SVD-LLM | 46.15 | 11.93 | |
| | SkipCat | 49.29 | 8.78 | |
| Qwen3-14B | - | Dense | 64.01 | - |
| | 20% | SVD-LLM | 55.11 | 8.89 |
| | | SkipCat | 60.15 | 3.85 |
| | 30% | SVD-LLM | 50.19 | 13.82 |
| SkipCat | | 56.04 | 7.97 | |

Table 2: Zero-shot accuracy of compressed larger models.

4.2 Performance Comparison

To assess overall performance, we compare *SkipCat* with existing low-rank compression methods in terms of perplexity and zero-shot task accuracy, as shown in Table 1. All approaches are evaluated without any additional training or the use of quantization techniques. *SkipCat* outperforms all other methods when evaluated under the same compression rate. On LLaMA2-7B with a 30% compression rate, our experiments show that existing low-rank compression methods fail to close the performance gap, increasing perplexity on WikiText-2 by 2.1× compared to the uncompressed model, whereas *SkipCat* incurs only a 1.4× increase. Notably, *SkipCat* at 30% compression achieves lower perplexity than all other methods compressed at just 20%. Moreover, in a training-free setting, our method achieves only a

| Comp. Rate | Ablation Settings | | | Perplexity ↓ | |
|------------|-------------------|------|--------|--------------|-------|
| | Cat | Skip | Quant. | Wiki2 | C4 |
| 20% | ✗ | ✗ | ✗ | 8.82 | 13.42 |
| | ✓ | ✗ | ✗ | 7.84 | 11.99 |
| | ✗ | ✓ | ✗ | 6.71 | 9.32 |
| | ✓ | ✓ | ✗ | 6.29 | 8.95 |
| | ✓ | ✓ | ✓ | 6.29 | 8.96 |

Table 3: Ablation study results on LLaMA2-7B.

2.2% average accuracy drop at 20% compression. When the compression rate increases to 30%, other methods exhibit a degradation of at least 13.54%, whereas *SkipCat* limits the drop to 6.34%, representing a relative improvement of 7%. The experiments on Qwen3-8B also show substantial improvements in both perplexity and zero-shot task accuracy, further confirming the effectiveness of our approach in compressing models.

To further assess the generalization capability of our method, we extend our experiments to larger models, including LLaMA2-13B and Qwen3-14B. As shown in Table 2, despite the increased model scale, *SkipCat* consistently achieves the most effective compression across all methods under the same compression rates.

4.3 Ablation Study

Our proposed method consists of two core techniques: *Cat* and *Skip*. These two components together form the foundation of *SkipCat*. To examine the individual effectiveness of each technique, we conduct an ablation study on LLaMA2-7B under a 20% compression rate. As shown in Table 3, using naive SVD low-rank compression alone increases the perplexity on WikiText-2 from 5.47 to 8.82. When only the *Cat* technique is applied, the perplexity decreases to 7.84. When only *Skip* is applied, it further reduces to 6.71. These results support our claim in Figure 1 that both techniques help increase the number of effective ranks, which in turn reduces performance degradation. When both *Cat* and *Skip* are applied together, the model achieves the lowest perplexity, demonstrating the complementary benefits of combining the two.

We evaluate the compatibility of our approach with quantization. Specifically, we apply HQQ (Badri and Shaji 2023) with rescaling to quantize the compressed model to 8-bit precision. The experimental results show that our parameter-level compression can be effectively combined with precision-level quantization. This enables even greater compression without introducing performance loss.

4.4 Improving Compression with Fine-Tuning

Our method already shows strong performance compared to existing low-rank compression approaches in a training-free setting. However, in many real-world applications such as domain-specific adaptation or deployment on edge devices, additional fine-tuning can offer further benefits. By combining low-rank compression with task-specific fine-tuning, it is possible not only to improve accuracy on downstream tasks,

but also to achieve higher compression rates. To investigate this potential, we fine-tune the low-rank compressed models at various compression levels. Following the setup in SVD-LLM, we apply LoRA-based fine-tuning (Hu et al. 2022) on the Alpaca (Taori et al. 2023) dataset to adapt the models.

In this experiment, we use LLaMA2-7B as the base model. The model is first compressed to a target compression ratio, followed by fine-tuning using LoRA. The hyperparameters are set according to the configuration used in SVD-LLM (Wang et al. 2025b). As shown in Table 4, we fine-tune both the *SkipCat*-compressed and SVD-LLM-compressed models and evaluate them on zero-shot tasks. At a 20% compression rate, we observe that after fine-tuning, the model compressed with *SkipCat* shows only a 0.39% drop in average accuracy compared to the dense model. In contrast, the model compressed with SVD-LLM still exhibits a 3.78% accuracy gap. This improvement can be attributed to our method’s ability to maximize the number of effectively utilized ranks through architectural design. By enabling richer capacity within the same compression budget, our approach allows the fine-tuned model to recover performance more effectively than naive low-rank methods. Even at more aggressive compression rates, our method maintains consistently higher accuracy than SVD-LLM.

| Comp. Rate | Method | Zero-shot Task Accuracy | |
|------------|----------------|-------------------------|--------------|
| | | Avg. (%) ↑ | Drop (%) ↓ |
| - | Dense | 54.79 | - |
| 20% | SVD-LLM | 51.02 | 3.78 |
| | <i>SkipCat</i> | 54.41 | 0.39 |
| 40% | SVD-LLM | 46.91 | 7.88 |
| | <i>SkipCat</i> | 48.65 | 6.15 |
| 60% | SVD-LLM | 39.50 | 15.29 |
| | <i>SkipCat</i> | 41.16 | 13.64 |
| 80% | SVD-LLM | 32.33 | 22.47 |
| | <i>SkipCat</i> | 32.82 | 21.97 |

Table 4: Zero-shot accuracy of fine-tuned compressed models based on LLaMA2-7B.

5 Conclusion

In this work, we propose *SkipCat*, a novel low-rank compression framework for large language models. Our method integrates two key techniques: *Cat* and *Skip*. The former allows multiple matrices within a single layer to share a common low-rank projection, while the latter, block skipping, reduces computation by omitting calculations in certain sub-blocks of the matrices. These components allow the model to increase the number of effectively utilized ranks without changing the overall compression ratio, which helps maximize performance. We conduct evaluations on downstream tasks and compare *SkipCat* with existing low-rank compression methods. Our method consistently achieves significantly higher accuracy under the same compression rates. These results demonstrate that increasing the number of effective ranks is a highly effective strategy for improving the quality of compressed models.

References

- Akhauri, Y.; AbouElhamayed, A. F.; Dotzel, J.; Zhang, Z.; Rush, A. M.; Huda, S.; and Abdelfattah, M. S. 2024. ShadowLLM: Predictor-based Contextual Sparsity for Large Language Models. *arXiv:2406.16635*.
- Amini, A.; Gabriel, S.; Lin, P.; Koncel-Kedziorski, R.; Choi, Y.; and Hajishirzi, H. 2019. MathQA: Towards Interpretable Math Word Problem Solving with Operation-Based Formalisms. *arXiv:1905.13319*.
- Badri, H.; and Shaji, A. 2023. Half-Quadratic Quantization of Large Machine Learning Models.
- Bisk, Y.; Zellers, R.; Bras, R. L.; Gao, J.; and Choi, Y. 2020. PIQA: Reasoning about Physical Commonsense in Natural Language. In *Thirty-Fourth AAAI Conference on Artificial Intelligence*.
- Clark, P.; Cowhey, I.; Etzioni, O.; Khot, T.; Sabharwal, A.; Schoenick, C.; and Tafjord, O. 2018. Think you have Solved Question Answering? Try ARC, the AI2 Reasoning Challenge. *ArXiv*, abs/1803.05457.
- Frantar, E.; and Alistarh, D. 2023. SparseGPT: Massive Language Models Can Be Accurately Pruned in One-Shot. *arXiv preprint arXiv:2301.00774*.
- Frantar, E.; Ashkboos, S.; Hoefler, T.; and Alistarh, D. 2022. GPTQ: Accurate Post-training Compression for Generative Pretrained Transformers. *arXiv preprint arXiv:2210.17323*.
- Gao, L.; Tow, J.; Biderman, S.; Black, S.; DiPofi, A.; Foster, C.; Golding, L.; Hsu, J.; McDonell, K.; Muennighoff, N.; et al. 2021. A framework for few-shot language model evaluation. *Version v0. 0.1. Sept*, 10: 8–9.
- Gu, M.; and Eisenstat, S. C. 1996. Efficient Algorithms for Computing a Strong Rank-Revealing QR Factorization. *SIAM Journal on Scientific Computing*.
- Hsu, Y.-C.; Hua, T.; Chang, S.; Lou, Q.; Shen, Y.; and Jin, H. 2022. Language model compression with weighted low-rank factorization. *arXiv preprint arXiv:2207.00112*.
- Hu, E. J.; Shen, Y.; Wallis, P.; Allen-Zhu, Z.; Li, Y.; Wang, S.; Wang, L.; Chen, W.; et al. 2022. Lora: Low-rank adaptation of large language models. *ICLR*, 1(2): 3.
- Kaplan, J.; McCandlish, S.; Henighan, T.; Brown, T. B.; Chess, B.; Child, R.; Gray, S.; Radford, A.; Wu, J.; and Amodei, D. 2020. Scaling Laws for Neural Language Models.
- Kaushal, A.; Vaidhya, T.; and Rish, I. 2023. Lord: Low rank decomposition of monolingual code llms for one-shot compression. *arXiv preprint arXiv:2309.14021*.
- Koike-Akino, T.; Chen, X.; Liu, J.; Wang, Y.; Brand, M.; et al. 2025. LatentLLM: Attention-Aware Joint Tensor Compression. *arXiv preprint arXiv:2505.18413*.
- Liu, Z.; Zhao, C.; Fedorov, I.; Soran, B.; Choudhary, D.; Krishnamoorthi, R.; Chandra, V.; Tian, Y.; and Blankevoort, T. 2024. SpinQuant-LLM quantization with learned rotations. *arXiv preprint arXiv:2405.16406*.
- Ma, X.; Fang, G.; and Wang, X. 2023. LLM-Pruner: On the Structural Pruning of Large Language Models. In *Advances in Neural Information Processing Systems*.
- Merity, S.; Xiong, C.; Bradbury, J.; and Socher, R. 2016. Pointer sentinel mixture models. *arXiv preprint arXiv:1609.07843*.
- Mihaylov, T.; Clark, P.; Khot, T.; and Sabharwal, A. 2018. Can a Suit of Armor Conduct Electricity? A New Dataset for Open Book Question Answering. In *EMNLP*.
- Qinsi, W.; Ke, J.; Tomizuka, M.; Keutzer, K.; and Xu, C. 2025. Dobi-SVD: Differentiable SVD for LLM Compression and Some New Perspectives. In *The Thirteenth International Conference on Learning Representations*.
- Raffel, C.; Shazeer, N.; Roberts, A.; Lee, K.; Narang, S.; Matena, M.; Zhou, Y.; Li, W.; and Liu, P. J. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of machine learning research*, 21(140): 1–67.
- Sakaguchi, K.; Bras, R. L.; Bhagavatula, C.; and Choi, Y. 2019. WinoGrande: An Adversarial Winograd Schema Challenge at Scale. *arXiv preprint arXiv:1907.10641*.
- Sun, M.; Liu, Z.; Bair, A.; and Kolter, J. Z. 2023. A Simple and Effective Pruning Approach for Large Language Models. *arXiv preprint arXiv:2306.11695*.
- Taori, R.; Gulrajani, I.; Zhang, T.; Dubois, Y.; Li, X.; Guestrin, C.; Liang, P.; and Hashimoto, T. B. 2023. Stanford alpaca: An instruction-following llama model.
- Touvron, H.; Martin, L.; Stone, K.; Albert, P.; Almahairi, A.; Babaei, Y.; Bashlykov, N.; Batra, S.; Bhargava, P.; Bhosale, S.; et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- Wang, J.; Chen, Y.-G.; Lin, I.-C.; Li, B.; and Zhang, G. L. 2024. Basis sharing: Cross-layer parameter sharing for large language model compression. *arXiv preprint arXiv:2410.03765*.
- Wang, X.; Alam, S.; Wan, Z.; Shen, H.; and Zhang, M. 2025a. Svd-llm v2: Optimizing singular value truncation for large language model compression. *arXiv preprint arXiv:2503.12340*.
- Wang, X.; Zheng, Y.; Wan, Z.; and Zhang, M. 2025b. SVD-LLM: Truncation-aware Singular Value Decomposition for Large Language Model Compression. In *International Conference on Learning Representations (ICLR)*.
- Xiao, G.; Lin, J.; Seznec, M.; Wu, H.; Demouth, J.; and Han, S. 2023. SmoothQuant: Accurate and Efficient Post-Training Quantization for Large Language Models. In *Proceedings of the 40th International Conference on Machine Learning*.
- Yang, A.; Li, A.; Yang, B.; Zhang, B.; Hui, B.; Zheng, B.; Yu, B.; Gao, C.; Huang, C.; Lv, C.; Zheng, C.; Liu, D.; Zhou, F.; Huang, F.; Hu, F.; Ge, H.; Wei, H.; Lin, H.; Tang, J.; Yang, J.; Tu, J.; Zhang, J.; Yang, J.; Yang, J.; Zhou, J.; Zhou, J.; Lin, J.; Dang, K.; Bao, K.; Yang, K.; Yu, L.; Deng, L.; Li, M.; Xue, M.; Li, M.; Zhang, P.; Wang, P.; Zhu, Q.; Men, R.; Gao, R.; Liu, S.; Luo, S.; Li, T.; Tang, T.; Yin, W.; Ren, X.; Wang, X.; Zhang, X.; Ren, X.; Fan, Y.; Su, Y.; Zhang, Y.; Zhang, Y.; Wan, Y.; Liu, Y.; Wang, Z.; Cui, Z.; Zhang, Z.; Zhou, Z.; and Qiu, Z. 2025. Qwen3 Technical Report. *arXiv preprint arXiv:2505.09388*.

Yuan, Z.; Shang, Y.; Song, Y.; Wu, Q.; Yan, Y.; and Sun, G. 2023. ASVD: Activation-aware Singular Value Decomposition for Compressing Large Language Models. arXiv:2312.05821.

Yuan, Z.; Shang, Y.; Zhou, Y.; Dong, Z.; Zhou, Z.; Xue, C.; Wu, B.; Li, Z.; Gu, Q.; Lee, Y. J.; et al. 2024. Llm inference unveiled: Survey and roofline model insights. *arXiv preprint arXiv:2402.16363*.

Zellers, R.; Holtzman, A.; Bisk, Y.; Farhadi, A.; and Choi, Y. 2019. HellaSwag: Can a Machine Really Finish Your Sentence? In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*.

Zhu, X.; Li, J.; Liu, Y.; Ma, C.; and Wang, W. 2024. A survey on model compression for large language models. *Transactions of the Association for Computational Linguistics*, 12: 1556–1577.