

FLAT-LLM: Fine-grained Low-rank Activation Space Transformation for Large Language Model Compression

Jiayi Tian¹, Ryan Solgi¹, Jinming Lu¹, Yifan Yang¹, Hai Li², Zheng Zhang¹

¹ University of California, Santa Barbara, ² Intel Corporation
{jiayi_tian, solgi, jinminglu, yifanyang}@ucsb.edu, hai.li@intel.com, zhengzhang@ece.ucsb.edu

Abstract

Large Language Models (LLMs) have enabled remarkable progress in natural language processing, yet their high computational and memory demands pose challenges for deployment in resource-constrained environments. Although recent low-rank decomposition methods offer a promising path for structural compression, they often suffer from accuracy degradation, expensive calibration procedures, and result in inefficient model architectures that hinder real-world inference speedups. In this paper, we propose FLAT-LLM, a fast and accurate, training-free structural compression method based on fine-grained low-rank transformations in the activation space. Specifically, we reduce the hidden dimension by transforming the weights using truncated eigenvectors computed via head-wise Principal Component Analysis, and employ a greedy budget redistribution strategy to adaptively allocate ranks across decoders. FLAT-LLM achieves efficient and effective weight compression without recovery fine-tuning, which could complete the calibration within a few minutes. Evaluated across 5 models and 11 datasets, FLAT-LLM outperforms structural pruning baselines in generalization and downstream performance, while delivering inference speedups over low-rank-based methods.¹

1 Introduction

Large Language Models (LLMs) have achieved state-of-the-art performance in a wide range of natural language processing and understanding tasks (Bai et al., 2023; Liu et al., 2024; Touvron et al., 2023). However, their substantial parameter counts and computational demands pose significant challenges for deployment in edge devices and resource-constrained environments. Model compression is a promising direction for reducing

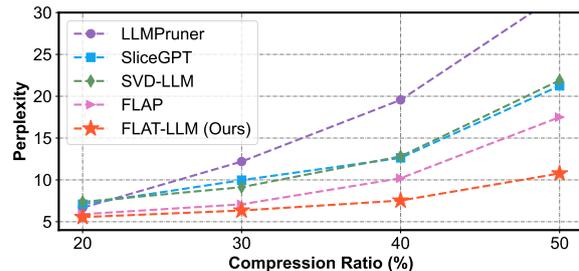


Figure 1: Comparison of WikiText-2 perplexity against various baselines on Llama-2 13B model.

both the model size and computational complexity, with popular approaches including quantization (Huang et al., 2024; Tian et al., 2023), knowledge distillation (Sun et al., 2020; Jiao et al., 2020), pruning (Sun et al., 2024; Ma et al., 2023; Yang et al., 2025), and low-rank decomposition (Ashkboos et al., 2024; Yang et al., 2024a). Among these, low-rank decomposition stands out for its hardware efficiency due to its structural nature.

Singular Value Decomposition (SVD) is a common low-rank decomposition approach for compressing the weight matrices of LLMs. However, weights from standard pre-training are often nearly full-rank and difficult to compress (Yu and Wu, 2023). To address this, recent work uses the low-rank nature of activation spaces by projecting weights into these subspaces to improve the compression ratio. For example, ASVD (Yuan et al., 2023) uses activation norms to transform the weights before decomposition, and SVD-LLM (Wang et al., 2024) enhances the accuracy by linking truncated singular values to reconstruction loss.

Despite recent improvements, SVD-based methods often incur performance drops and require recovery fine-tuning, even under low compression ratios. This stems from a key limitation: preserving sufficient information via high-rank SVD increases model parameters, as both left and right singular vectors must be stored. The problem is especially

¹Code is available in <https://github.com/TTTTTTris/FLAT-LLM>.

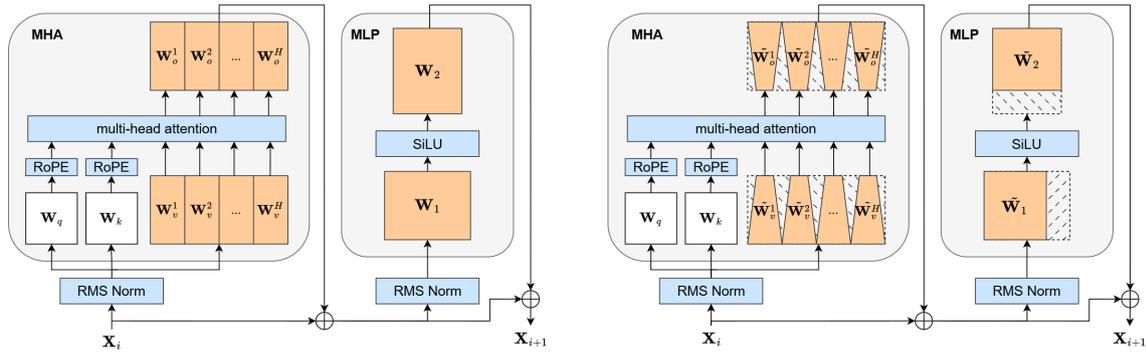


Figure 2: Decoder structure before (left) and after (right) weight truncation. Orange blocks indicate truncated weights; hatched areas show removed weights; blue boxes denote non-linear functions.

severe for square matrices (which are common in LLMs like Llama), where reducing parameters requires truncating at least 50% of singular values, often leading to significant information loss. Moreover, replacing a single large matrix multiplication with two smaller ones in SVD-format linear layers can degrade GPU efficiency, limiting practical inference speedups.

To address the inherent limitations of SVD, SliceGPT (Ashkboos et al., 2024) projects post-RMSNorm hidden states into low-rank subspaces using Principal Component Analysis (PCA), enabling the resulting eigenvector pairs to be fully absorbed into adjacent layers. However, to enable this transformation between decoder layers, SliceGPT requires inserting adapter modules along the residual paths, which bring high memory overhead and limit the inference speedup. For example, on Llama-2 7B, SliceGPT incur 10% additional memory overhead at a 20% compression ratio, and yields only $1.1\times$ speedup in inference at a 50% compression ratio.

To further enhance the performance of low-rank compressed models, prior works have introduced rank-adaptive methods that allocate heterogeneous ranks across different model components. For example, Adaptive SVD (Gao et al., 2024b) employs a trainable hypernetwork to determine layer-wise ranks, achieving higher compression efficiency than uniform baselines but at the cost of task-specific retraining. However, this approach requires weeks on LLaMA-7B, making it impractical and highly inefficient for scaling to larger models such as LLaMA-70B. MoDeGPT (Lin et al., 2024) adopts an importance-score-based sparsity allocation scheme with entropic regularization for smoothing, but it demands extensive hyperparam-

eter tuning to obtain competitive performance.

To overcome the challenges mentioned above, we propose FLAT-LLM, a fast, accurate, and training-free structural compression method for LLMs. Figure 2 illustrates the decoder architecture before and after applying FLAT-LLM compression. FLAT-LLM projects the post-value hidden states into low-rank subspaces using fine-grained, head-wise PCA. The resulting low-rank eigenvector pairs are then absorbed into the value and output weight matrices to complete the compression. Unlike previous approaches, this joint absorption-based compression introduces no additional memory overhead, and the overall compression ratio directly corresponds to the retained rank ratio. To further improve performance, we introduce an importance-preserving rank selection algorithm. This algorithm is entirely tuning-free and completes within seconds, achieving over $100\times$ higher time efficiency compared to Adaptive SVD (Gao et al., 2024b), and consistently outperforming MoDeGPT’s rank selection strategy. Our main contributions are as follows:

- We propose a training-free, fine-grained compression technique that operates within multi-head attention layers, avoiding the inefficiencies of prior decomposition-based methods.
- We introduce a novel training-free rank selection algorithm that allocates ranks using a greedy redistribution strategy and can be integrated with existing low-rank LLM compression pipelines.
- We demonstrate the effectiveness of FLAT-LLM through extensive evaluations on language modeling and downstream tasks. As shown in Figure 1, FLAT-LLM significantly improves perplexity on the WikiText-2 test split across a range

of compression ratios, indicating enhanced text generation capabilities under varying levels of compression.

2 Related Work

In addition to structural pruning via weight decomposition (as discussed in the introduction), another line of research achieves structural compression by directly removing model components based on importance scores. These approaches can be broadly categorized into two types: fine-grained pruning and coarse-grained pruning. The former offers fine-grained control by removing individual rows or columns within weight matrices. For example, LLM-Pruner (Ma et al., 2023) leverages gradient-based saliency scores to prune less important components, while FLAP (An et al., 2024) adaptively removes unstable neurons and channels based on activation fluctuations.

In contrast, the coarse-grained pruning eliminates larger components such as attention heads, layers, or entire decoders. Though more efficient for inference, this often results in greater performance degradation. For instance, ShortGPT (Men et al., 2024) prunes decoders based on cosine similarity-based importance ranking; LaCo (Yang et al., 2024b) compresses models by merging adjacent layers; and BlockPruner (Zhong et al., 2024) removes redundant MHA or MLP blocks through iterative search. In this work, we demonstrate that our method outperforms both fine- and coarse-grained structural pruning baselines in terms of performance under compression.

3 Background

Multi-Head Attention. A typical LLM consists of multiple decoder layers, each comprising two main components: Multi-Head Attention (MHA) and a Multi-Layer Perceptron (MLP). We define the input hidden states $\mathbf{X} \in \mathbb{R}^{N \times d_{hid}}$, where N denote the sequence length, and d_{hid}, d_h represent the dimension of hidden states and each attention head, respectively. Here, the MHA module computes a weighted aggregation of values using learned projections of queries, keys, and values. For a total of H attention heads, each head h performs:

$$\begin{aligned} \mathbf{A}^h &= \frac{\mathbf{X}\mathbf{W}_q^{h\top}(\mathbf{X}\mathbf{W}_k^{h\top})^\top}{\sqrt{d_h}}, \\ \mathbf{Y}_v^h &= \text{Softmax}(\mathbf{A}^h)\mathbf{X}\mathbf{W}_v^{h\top}, \end{aligned} \quad (1)$$

where $\mathbf{W}_q^h, \mathbf{W}_k^h, \mathbf{W}_v^h \in \mathbb{R}^{d_h \times d_{hid}}$ are the projection matrices in query, key and value layer in head h . The attention matrix $\mathbf{A}^h \in \mathbb{R}^{N \times N}$ captures token-wise interactions, and the Softmax function is applied to compute attention scores, which are then used to weight the value representations. The result $\mathbf{Y}_v^h \in \mathbb{R}^{N \times d_h}$ represents the per-head value output. This is further transformed by a learned output projection $\mathbf{W}_o^h \in \mathbb{R}^{d_{hid} \times d_h}$, yielding the partial attention output $\mathbf{Y}_o^h \in \mathbb{R}^{N \times d_{hid}}$. The final output of the multi-head attention layer is then obtained by aggregating the partial outputs from all heads:

$$\begin{aligned} \mathbf{Y}_o^h &= \text{Softmax}(\mathbf{A}^h)\mathbf{X}\mathbf{W}_v^{h\top}\mathbf{W}_o^{h\top}, \\ \mathbf{Y}_o &= \text{sum}(\mathbf{Y}_o^1, \dots, \mathbf{Y}_o^H). \end{aligned} \quad (2)$$

Principle Component Analysis (PCA). Principal Component Analysis (PCA) is a classical dimensionality reduction technique that identifies the principal components along which the data exhibit the greatest variance. Given a data matrix $\mathbf{Z} \in \mathbb{R}^{N \times d}$, PCA computes the eigen decomposition of the covariance matrix \mathbf{C} :

$$\mathbf{C} = \mathbf{Z}^\top \mathbf{Z} = \mathbf{Q}\mathbf{\Lambda}\mathbf{Q}^\top, \quad (3)$$

where $\mathbf{Q} \in \mathbb{R}^{d \times d}$ is the orthogonal matrix of eigenvectors, and $\mathbf{\Lambda} \in \mathbb{R}^{d \times d}$ is the diagonal matrix of corresponding eigenvalues. To capture the most significant variance, we retain the top- r principal components and define $\tilde{\mathbf{Q}} \in \mathbb{R}^{d \times r}$ as the truncated eigenvector matrix. The corresponding rank- r reconstruction of the original data is defined as projecting \mathbf{Z} onto the rank- r space and then mapping it back to the original space:

$$\tilde{\mathbf{Z}} = \mathbf{Z}\tilde{\mathbf{Q}}\tilde{\mathbf{Q}}^\top. \quad (4)$$

PCA thus reveals low-rank structure in weight matrices or hidden representations, enabling efficient approximations with minimal loss of information.

4 FLAT-LLM

In this section, we present FLAT-LLM in detail, beginning with a head-wise PCA-based weight truncation method that compresses weight matrices by truncating and absorbing the eigenvectors. We then introduce an importance-preserving rank selection strategy to allocate adaptive ranks across decoder layers. Finally, we conduct a theoretical analysis of the truncation loss in our head-wise PCA method.

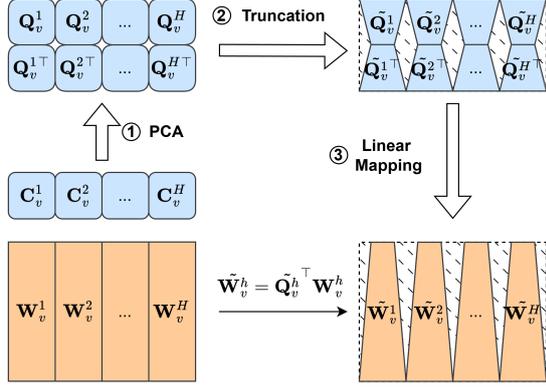


Figure 3: Fine-grained head-wise PCA in value layer.

4.1 Head-wise PCA-based Weight Truncation

Inspired by Equation (2), we observe that the value and output projections within each attention head are computed consecutively. Leveraging this observation, we propose to exploit the low-rank activation space of the value output to jointly compress the value and output weight matrices.

As shown in Figure 3, the detail compression process with M calibration samples are given in three steps: ① Compute the covariance matrix $\mathbf{C}_v^h = \sum_{m=1}^M \mathbf{Y}_{v,m}^h \mathbf{Y}_{v,m}^{h\top}$ and perform PCA to obtain the orthogonal eigenvector $\mathbf{Q}_v^h \in \mathbb{R}^{d_h \times d_h}$ utilizing Equation (3). ② Truncate the eigenvectors to rank r , yielding reduced basis $\tilde{\mathbf{Q}}_v^h \in \mathbb{R}^{d_h \times r}$, and the reconstructed per-head value output becomes $\tilde{\mathbf{Y}}_v^h = \mathbf{Y}_v^h \tilde{\mathbf{Q}}_v^h \tilde{\mathbf{Q}}_v^{h\top}$. ③ To compress the weights by absorbing the truncated eigenvectors, we reformulate the MHA computation in Equation (2) as the following:

$$\mathbf{Y}_o^h = \text{Softmax}(\mathbf{A}^h) \mathbf{X} \mathbf{W}_v^{h\top} \mathbf{Q}_v^h \mathbf{Q}_v^{h\top} \mathbf{W}_o^{h\top} \quad (5)$$

This enables the jointly compression on the value and output weights using the PCA basis derived from the value layer:

$$\begin{aligned} \tilde{\mathbf{Y}}_o^h &= \text{Softmax}(\mathbf{A}^h) \mathbf{X} \mathbf{W}_v^{h\top} \tilde{\mathbf{Q}}_v^h \tilde{\mathbf{Q}}_v^{h\top} \mathbf{W}_o^{h\top}, \\ \tilde{\mathbf{Y}}_o^h &= \text{Softmax}(\mathbf{A}^h) \mathbf{X} \tilde{\mathbf{W}}_v^{h\top} \tilde{\mathbf{W}}_o^{h\top}, \end{aligned} \quad (6)$$

where the first and second equations represent the truncation and absorption. Here, we aim for $\tilde{\mathbf{Q}}_v^h \tilde{\mathbf{Q}}_v^{h\top} \approx \mathbf{I}_v^h$ to retain most of the representational power of the original \mathbf{Y}_v^h . After absorbing the truncated basis into the weights, the value and output projections of each head are reduced to $\tilde{\mathbf{W}}_v^h \in \mathbb{R}^{r \times d_{hid}}$, $\tilde{\mathbf{W}}_o^h \in \mathbb{R}^{d_{hid} \times r}$, respectively.

In this way, we can jointly compress both the value and output weights to $\frac{r}{d_h}$ of their original

size leveraging the low-rank structure of the output hidden states from the value layer. Notably, this joint compression technique remains compatible with modern architectures using Grouped-Query Attention (GQA) (Ainslie et al., 2023), resulting in different numbers of value and output heads. The detailed formulation for the GQA case is provided in Appendix A.

Additionally, although the query and key projections cannot be jointly compressed in the same manner as the value and output weight matrices, head-wise PCA can still be applied to them independently. Specifically, we follow the same three-step head-wise PCA procedure illustrated in Figure 3. The reconstructed query and key outputs for head h can be written as

$$\tilde{\mathbf{Y}}_q^h = \mathbf{Y}_q^h \tilde{\mathbf{Q}}_q^h \tilde{\mathbf{Q}}_q^{h\top}, \tilde{\mathbf{Y}}_k^h = \mathbf{Y}_k^h \tilde{\mathbf{Q}}_k^h \tilde{\mathbf{Q}}_k^{h\top}. \quad (7)$$

Accordingly, the compressed query and key projection matrices are given by

$$\tilde{\mathbf{W}}_q^h = \tilde{\mathbf{Q}}_q^h \mathbf{W}_q^{h\top}, \quad \tilde{\mathbf{W}}_k^h = \tilde{\mathbf{Q}}_k^h \mathbf{W}_k^{h\top}. \quad (8)$$

Since the reduced basis $\tilde{\mathbf{Q}}_v^h \in \mathbb{R}^{d_h \times r}$ is $H \times$ is significantly smaller than the corresponding compressed projection matrices $\tilde{\mathbf{W}}_{q,k}^h \in \mathbb{R}^{r \times d_{hid}}$, performing head-wise PCA on the query and key projections leads to substantially lower approximation error than applying a full-matrix low-rank decomposition at the same sparsity level.

4.2 Importance-Preserving Rank Selection

In our experiments, we observed that using a uniform rank across all layers degrades performance, particularly under high compression ratios. We first analyze the cosine similarity between each decoder layer’s input and output hidden states, and reveal that the intrinsic dimensionality varies across layers. Motivated by this observation, we propose a decoder-wise rank selection algorithm that employs a greedy redistribution strategy to adaptively allocate ranks based on their relative importance.

To analyze variations in intrinsic dimension across decoder layers, we compute the cosine similarity between the input and output hidden states of each decoder. Given a model with L decoder layers, let \mathbf{X}_l and \mathbf{X}_{l+1} denote the input and output hidden state matrices of the l -th decoder layer, respectively. For each sample $p \in \{1, \dots, N\}$, the cosine similarity $c_{l,p}$ between the corresponding

Algorithm 1 Importance-Preserving Rank Selection

Require: Sparsity s , number of decoders L , importance scores \mathbf{t}

- 1: Initialize total budget $B \leftarrow L(1 - s)$, active set $\mathcal{A} \leftarrow \{1, \dots, L\}$
 - 2: **while** $\mathcal{A} \neq \emptyset$ **do**
 - 3: $\tilde{w}_l = \frac{t_l}{\sum_{j \in \mathcal{A}} t_j} \cdot B$ for all $l \in \mathcal{A}$
 - 4: Let $\mathcal{S} \leftarrow \{l \in \mathcal{A} \mid \tilde{w}_l > 1\}$
 - 5: **if** $\mathcal{S} = \emptyset$ **then** \triangleright Assign all remain entries
 - 6: $w_l \leftarrow \tilde{w}_l$ for all $l \in \mathcal{A}$
 - 7: **break**
 - 8: **end if**
 - 9: **for all** $l \in \mathcal{S}$ **do** \triangleright Assign fixed entries
 - 10: $w_l \leftarrow 1$
 - 11: $B \leftarrow B - w_l$
 - 12: **end for**
 - 13: $\mathcal{A} \leftarrow \mathcal{A} \setminus \mathcal{S}$ \triangleright Remove fixed entries
 - 14: **end while**
 - 15: **return** \mathbf{w} \triangleright Final allocation
-

rows of \mathbf{X}_l and \mathbf{X}_{l+1} is defined as:

$$c_{l,p} = \frac{\mathbf{X}_{l,p}^\top \mathbf{X}_{l+1,p}}{\|\mathbf{X}_{l,p}\|_2 \|\mathbf{X}_{l+1,p}\|_2}, \quad (9)$$

where $\mathbf{X}_{l,p}$ and $\mathbf{X}_{l+1,p}$ represent the p -th row of \mathbf{X}_l and \mathbf{X}_{l+1} , respectively. The average cosine similarity for the l -th decoder layer is then given by $c_l = \mathbb{E}_{X,p}[c_{l,p}]$, reflecting the overall alignment between input and output hidden states.

To quantify the degree of dissimilarity and thereby infer the intrinsic dimension, we compute $t_l = \frac{\arccos c_l}{\pi}$, which captures the normalized angular deviation between the representations. We interpret t_l as an indicator of the relative importance and compressibility of each decoder layer. From empirical evaluations on several LLMs, we observe that t_l typically ranges from 0.06 to 0.30. This leads to two key insights: (1) intrinsic dimensionality varies across layers, motivating the use of heterogeneous rank assignments; and (2) the small angular deviations suggest that decoder layers often reside in low-dimensional subspaces, indicating that LLMs are amenable to compression.

Given the model sparsity as s , to determine the remaining rank ratios w_l for l -th decoder regarding the importance score t_l , a naive way is to proportional scale the total budget of remaining rank ratios $B = L(1 - s)$, which gives $\hat{w}_l = \frac{t_l}{\sum_{l=1}^L t_l} B$. However, this naive proportional allocation may

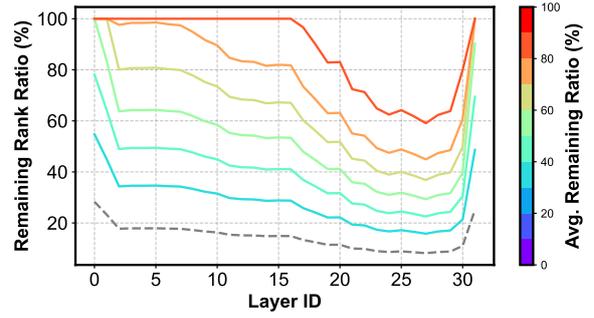


Figure 4: Remaining rank ratio versus layer id computed with Algorithm 1. The average remaining ratio is set between 30% (lowest solid) to 90% (highest solid).

violate the $w_l \in [0, 1]$ constraint when some components receive disproportionately high scores. In order to design w_l regarding the importance score t_l , while fit each w_l within the constraint and fix the total budget for remaining size as $\sum_l w_l = B$, the objective of our rank selection algorithm can be defined as:

$$\min_{\mathbf{w} \in [0,1]^L} \|\mathbf{w} - \hat{\mathbf{w}}\|, \text{ s.t. } \sum_l w_l = B, \quad (10)$$

where $\mathbf{w} = [w_1, \dots, w_L]$ and $\hat{\mathbf{w}} = [\hat{w}_1, \dots, \hat{w}_L]$ are vectors constructed by the ratios w_l and \hat{w}_l in each decoders.

To address this, we implement a greedy redistribution strategy, as shown in Algorithm 1. First, we define a variable \tilde{w}_l to represent the proportional scaled ratios \hat{w}_l when the budget B is changing during the rank selection process. Given an active set \mathcal{A} that contains the indices of unassigned remaining rank ratios, we iteratively compute \tilde{w}_l in the active set with latest budget B , clip \tilde{w}_l that exceed the upper bound as 1, update the total budget B , and remove the clipped entries from the active set. This process continues until all elements has been assigned a remaining rank ratios w_l . In this way, the resulting solution w_l remains proportional to t_l where possible, while ensuring the boundedness and budget feasibility.

Figure 4 illustrates the remaining rank ratios across the decoder blocks of Llama-2 7B under different average remaining ratios (from 40% to 90%) determined by the IPRS algorithm. The gray dashed curve represents the normalized importance scores \mathbf{t} that guide the optimization process, while the colorful solid lines show the remaining rank ratio \mathbf{w} under multiple sparsity ratio $s \in [0.1, 0.6]$. As shown, the resulting rank allocations are non-uniform and adaptively preserve more capacity in

layers with higher importance. Additional visualizations of the rank distributions produced by IPRS on other models are provided in Appendix D.4.

4.3 Truncation Error Analysis

In the following, we provide the theoretical proof of the direct correspondence between eigenvalues and truncation loss in both single-head and multi-head cases, which guarantees the effectiveness of our method. The detailed derivations are presented in Appendix B.1. We also include an empirical comparison of reconstruction errors in Appendix B.2 with SVD-LLM to further demonstrate the advantage of our head-wise PCA for the attention block.

As described in Section 4.1, to compress the value layer, we perform PCA on its feature covariance and project onto a low-rank subspace:

$$\tilde{\mathbf{Y}}_v^h = \mathbf{Y}_v^h \tilde{\mathbf{Q}}_v^h \tilde{\mathbf{Q}}_v^{h\top}, \quad (11)$$

where $\mathbf{Y}_v^{h\top} \mathbf{Y}_v^h = \mathbf{Q}_v^h \mathbf{\Lambda}_v^h \mathbf{Q}_v^{h\top}$, $\tilde{\mathbf{Q}}_v^h = \mathbf{Q}_v^h[:, :r]$.

Here, $\mathbf{Q}_v^h \in \mathbb{R}^{d_h \times d_h}$ contains orthonormal eigenvectors and $\mathbf{\Lambda}_v^h$ is diagonal with eigenvalues $\lambda_1^h \geq \dots \geq \lambda_{d_h}^h \geq 0$, and $\tilde{\mathbf{Q}}_v^h \in \mathbb{R}^{d_h \times r}$ contains the top- r principal components.

Theorem 4.1 (Reconstruction Error of Single-head Output PCA Projection). *Let $\mathbf{Y}_v^h = \mathbf{XW}_v^{h\top}$, and $\tilde{\mathbf{Y}}_v^h = \mathbf{Y}_v^h \tilde{\mathbf{Q}}_v^h \tilde{\mathbf{Q}}_v^{h\top}$ be the rank- r approximation obtained by projecting \mathbf{Y}_v^h onto its top- r principal components. Then the squared Frobenius norm of the reconstruction error satisfies:*

$$\|\mathbf{Y}_v^h - \tilde{\mathbf{Y}}_v^h\|_F^2 = \sum_{i=r+1}^{d_h} \lambda_i^h,$$

where $\{\lambda_i^h\}$ are the eigenvalues of $\mathbf{Y}_v^{h\top} \mathbf{Y}_v^h$.

Corollary 4.2 (Reconstruction Error of Multi-head Output PCA Projection). *Let $\mathbf{Y}_v = \text{concat}(\mathbf{Y}_v^1, \dots, \mathbf{Y}_v^H)$ be the concatenated output hidden states. The squared Frobenius norm of the reconstruction error satisfies:*

$$\|\mathbf{Y}_v - \tilde{\mathbf{Y}}_v\|_F^2 = \sum_{h=1}^H \sum_{i=r+1}^{d_h} \lambda_i^h,$$

Therefore, the reconstruction error of the multi-head value output can be expressed as the sum of the truncated eigenvalues λ_i^h from the output of each head value projection \mathbf{Y}_v^h . When the preserved dimension is r , truncating the smallest eigenvalues $\lambda_{r+1}^h, \dots, \lambda_{d_h}^h$ in the PCA decomposition of each value head yields the minimal reconstruction error for the multi-head value output.

5 Experiments

In this section, we first compare the performance of FLAT-LLM on language modeling and downstream tasks with recent fine-grained structural pruning methods across multiple LLM architectures and compression ratios. We then evaluate inference speedup and memory saving, and conduct ablation studies on our IPRS algorithm and the impact of calibration datasets. Additional experimental results, including calibration efficiency, performance under varying compression ratios on downstream tasks, and comparisons with coarse-grained importance-based pruning baselines, are provided in Appendix D.

5.1 Setups

Models and Datasets. We evaluate our method on multiple decoder-based generative models, including Llama-2 7B, 13B, 70B, Llama-3 8B (Touvron et al., 2023), and Mistral 7B-v0.1 (Jiang et al., 2023). Following the settings in previous works (Ashkboos et al., 2024; Wang et al., 2024), we use 256 samples with 4096 tokens from WikiText-2 or Alpaca datasets for calibration. For the downstream evaluation, we use the LM Evaluation Harness (Gao et al., 2024a) and test on ten tasks: ARC-e, ARC-c, PIQA, WinoGrande, HellaSwag, BoolQ, OBQA, MathQA, CommonsenseQA, MMLU. Here, MMLU uses 5-shot evaluation, and all others use zero-shot.

Baselines. We compare our method with recent decomposition-based pruning approaches, including SVD-LLM (Wang et al., 2024), SliceGPT (Ashkboos et al., 2024), as well as fine-grained importance-based methods such as FLAP (An et al., 2024) and LLM-Pruner (Ma et al., 2023). Additionally, we evaluate our approach against coarse-grained structural pruning techniques, including Pivoting Factorization (Zhao et al., 2025), LLM Surgeon (van der Ouderaa et al., 2024), ShortGPT (Men et al., 2024), LaCo (Yang et al., 2024b), and BlockPruner (Zhong et al., 2024), as reported in Appendix D.2 and D.3.

Implementation Details. Our implementation is built on the Huggingface Transformers library (Wolf et al., 2019). The MHA blocks are compressed using our proposed head-wise PCA, while the MLP modules are compressed using selection matrices derived from the Nyström approximation; implementation details are provided in Appendix C.

Table 1: Comparison of downstream performance against prior structural compression methods on LLaMA-2 (7B, 13B, 70B), LLaMA-3 8B, and Mistral-7B models at 20% compression ratio. † indicates FLAT-LLM with head-wise query and key compression.

Model	Method	PPL ↓	Avg. ↑	MMLU 5-shot	PIQA	WinoG.	HellaS.	ARC-e	ARC-c	OBQA
LLaMA-2 7B	Original	5.11	62.14	45.70	79.05	69.38	75.92	74.49	46.25	44.20
	LLM-Pruner (Ma et al., 2023)	10.55	53.89	26.20	75.95	63.38	67.83	64.31	39.93	39.60
	FLAP (An et al., 2024)	6.76	53.07	31.90	74.54	62.98	64.74	61.28	36.43	39.60
	SliceGPT (Ashkboos et al., 2024)	8.24	46.26	26.75	64.80	62.98	49.18	55.68	31.40	33.00
	SVD-LLM (Wang et al., 2024)	7.84	52.32	29.34	71.49	65.27	58.57	68.31	35.84	37.40
	FLAT-LLM (ours) FLAT-LLM† (ours)	6.70 6.34	55.16 57.85	39.67 43.99	72.20 73.18	65.82 68.35	64.72 68.99	64.44 66.79	38.65 40.87	40.60 42.80
LLaMA-2 13B	Original	4.57	65.70	55.40	80.41	72.53	79.41	77.39	49.15	45.60
	LLM-Pruner (Ma et al., 2023)	9.67	55.45	22.80	77.97	60.77	71.26	67.09	44.28	44.00
	FLAP (An et al., 2024)	5.90	57.00	41.20	75.57	67.25	69.19	65.91	39.08	40.80
	SliceGPT (Ashkboos et al., 2024)	7.10	50.58	35.49	65.18	65.67	52.30	59.26	36.77	39.40
	SVD-LLM (Wang et al., 2024)	7.37	55.86	35.54	72.91	67.17	63.47	71.00	39.93	41.00
	FLAT-LLM (ours)	5.55	63.00	54.72	75.84	72.06	73.36	75.59	46.25	43.20
LLaMA-2 70B	Original	3.12	71.38	68.80	82.75	77.82	83.80	80.72	57.17	48.60
	FLAP (An et al., 2024)	8.76	48.29	25.90	72.31	64.09	55.94	51.05	31.91	36.80
	SliceGPT (Ashkboos et al., 2024)	5.76	57.40	48.30	68.01	72.14	57.16	68.64	43.94	43.60
	SVD-LLM (Wang et al., 2024)	5.96	61.07	52.10	74.48	72.61	68.41	71.93	46.93	41.00
	FLAT-LLM (ours)	4.33	67.98	67.35	77.20	77.03	78.44	78.87	51.19	45.80
	Mistral-7B	Original	4.92	68.14	62.50	82.05	73.95	81.02	79.55	53.92
FLAP (An et al., 2024)		7.11	48.29	25.90	72.31	64.09	55.94	51.05	31.91	36.80
SliceGPT (Ashkboos et al., 2024)		9.06	43.18	25.52	59.35	61.21	45.11	51.60	30.29	29.20
SVD-LLM (Wang et al., 2024)		9.29	50.23	25.02	70.46	64.56	58.09	69.28	37.63	26.60
FLAT-LLM (ours)		6.11	58.54	57.21	70.84	69.30	63.76	68.73	41.72	38.20
LLaMA-2 7B		Original	5.75	68.14	65.43	80.85	73.40	79.17	80.09	53.24
	FLAP (An et al., 2024)	8.42	54.42	42.24	73.50	65.90	59.87	64.22	36.26	39.00
	SliceGPT (Ashkboos et al., 2024)	16.62	41.28	25.07	60.23	57.22	40.46	47.26	27.73	31.00
	SVD-LLM (Wang et al., 2024)	17.17	47.06	28.64	66.27	61.01	52.65	55.43	31.66	33.80
	FLAT-LLM (ours)	8.15	61.47	62.56	72.80	72.53	69.13	68.90	42.58	41.80
	FLAT-LLM† (ours)	8.02	62.45	63.21	73.67	74.11	70.46	70.24	44.28	41.20

Unless otherwise specified, all main-text experimental results are obtained using FLAT-LLM with compression applied only to the value and output projections in the attention module. We denote the variant that additionally compresses the query and key projections as FLAT-LLM[†]. All decomposition computations are performed in double precision to ensure numerical stability. Experiments are conducted on a single A100 40GB GPU, except for LLaMA-2 70B, which is evaluated using 4 A100 GPUs. All methods are evaluated without any recovery fine-tuning.

5.2 Performance Comparison

Performance on Different LLMs. Table 1 compares the language modeling perplexity and downstream task accuracy of FLAT-LLM and FLAT-LLM[†] against prior structural compression methods across five LLMs under a 20% compression ratio. FLAT-LLM[†] augments FLAT-LLM by additionally applying head-wise compression to the query and key projections, which consistently improves accuracy under the same compression budget. Across all evaluated models, FLAT-LLM^(†) achieves the best overall performance in both av-

erage accuracy and perplexity, indicating strong generalization and text generation quality under compression. For larger models, like LLaMA-2 13B and 70B, FLAT-LLM incurs only a modest average accuracy drop of 2.7% and 3.4%. In contrast to prior methods, which often struggle on newer architectures like Mistral-7B and LLaMA-3 8B, FLAT-LLM consistently maintains high performance, yielding up to 15.4% and 20.2% accuracy improvements, respectively. Notably, FLAT-LLM outperforms all baselines by as much as 40% on the MMLU 5-shot benchmark, demonstrating its effectiveness at preserving factual reasoning and broad-domain knowledge.

We further evaluate zero-shot downstream performance under varying compression ratios (see Figure 8 in Appendix D.3). FLAT-LLM consistently outperforms prior low-rank decomposition baselines across all models and compression levels, demonstrating strong robustness and effectiveness.

5.3 Comparison with Quantization

We further integrate FLAT-LLM with post-training quantization methods to assess their combined effectiveness. As shown in Table 2, applying FLAT-

Table 2: Perplexity comparison with GPTQ for LLaMA-2 7B and Llama-3 8B on WikiText-2.

Model	Method	Size (GB)	PPL ↓
LLaMA-2 7B	Original	14	5.12
	GPTQ 2-bit	1.8 (7.8 ×)	NaN
	FLAT-LLM + GPTQ 3-bit	1.8 (7.8 ×)	13.43
LLaMA-3 8B	Original	16	5.75
	GPTQ 3-bit	3.0 (5.3 ×)	39.85
	FLAT-LLM + GPTQ 4-bit	2.8 (5.7 ×)	11.91

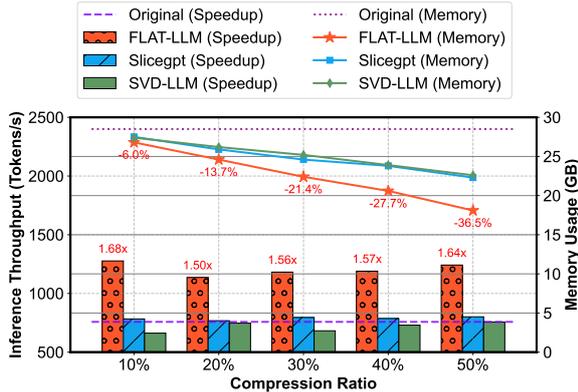


Figure 5: Comparison of inference throughput and memory usage with prior low-rank-based methods.

LLM (with a 30% compression ratio) followed by GPTQ-3/4-bit quantization achieves lower perplexity than GPTQ-2/3-bit alone, while maintaining the same or even lower memory footprint. This integration makes it possible to compress large language models by 5 – 8× with negligible accuracy degradation, enabling highly efficient and effective model deployment.

5.4 Inference Efficiency

Figure 5 compares the inference throughput and memory usage of our method with prior low-rank approaches, including SliceGPT and SVD-LLM, across compression ratios ranging from 10% to 50% on LLaMA-2 7B. Following previous setups (Ashkboos et al., 2024), we generate 256 tokens with a batch size of 64, where we report the throughput during decoding and the CUDA memory usage after decoding. Without any CUDA-level optimization, our method consistently outperforms the original model, achieving speedups over 1.50× across all compression ratios. Notably, even at a modest 10% overall compression ratio, our method attains a 1.68× throughput improvement, as the proposed non-uniform rank selection strategy concentrates compression on middle-to-late layers,

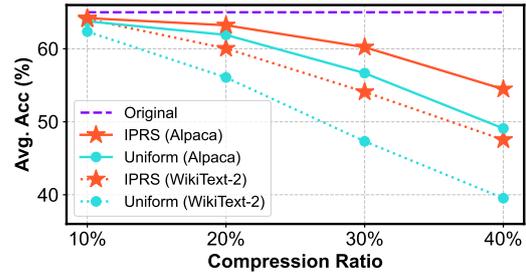


Figure 6: Comparison of zero-shot average accuracy on downstream datasets versus compression ratio using uniform rank and our IPRS algorithm on Llama-2 13B.

which are empirically more memory-bound in inference. Additionally, due to the reduced dimensionality of the value outputs, FLAT-LLM further reduces activation memory in the value cache. As a result, even under the same compression ratio, it achieves the lowest memory usage among all low-rank model compression methods. Compared to SliceGPT and SVD-LLM, it achieves up to 1.63× and 1.93× higher throughput, and reduces memory usage by 19% and 20%, respectively. These results highlight the efficiency of FLAT-LLM, making it a strong candidate for real-world deployment.

5.5 Ablation Study

In this section, we present an ablation study of our rank selection method and evaluate its performance using different calibration datasets. Specifically, we report the average accuracy across eight downstream tasks using calibration performed on either the WikiText-2 or Alpaca dataset.

Importance-Preserving Rank Selection. To evaluate the effectiveness of our Importance-Preserving Rank Selection (IPRS) method, Figure 6 compares the average zero-shot precision of different rank allocation strategies - uniform and IPRS - between compression ratios ranging from 10% to 40% on LLaMA-2 13B. As shown in Figure 6, IPRS consistently outperforms the uniform baseline across all compression settings. The performance gap becomes more pronounced at higher compression ratios, achieving gains of 5.4% with Alpaca calibration and 8.0% with WikiText-2 at a 40% compression ratio. This is because sensitive layers suffer disproportionately high truncation loss under aggressive compression, making adaptive rank selection increasingly important. These results underscore the complementary strengths of fine-grained PCA and importance-aware rank allo-

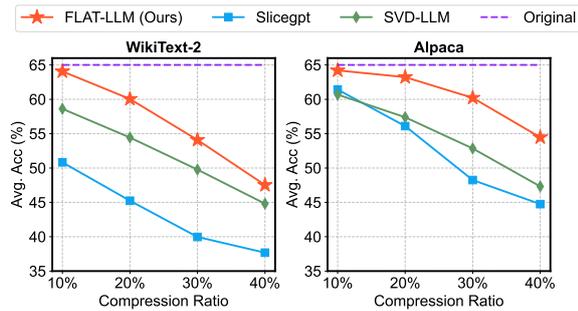


Figure 7: Comparison of zero-shot average accuracy on downstream datasets versus compression ratio calibrated on WikiText-2 or Alpaca on Llama-2 13B.

cation. In Appendix D.5, we further apply IPRS to the SVD-LLM method, which also shows improved performance over uniform rank. In Appendix D.6, we additionally show that IPRS consistently outperforms MoDeGPT’s rank allocation strategy.

Calibration Dataset. We evaluate the average zero-shot accuracy on LLaMA-2 13B using two calibration datasets—WikiText-2 and Alpaca. As shown in Figure 7, FLAT-LLM consistently outperforms SliceGPT and SVD-LLM across all settings and compression levels, demonstrating strong generalization with WikiText-2 calibration and even greater gains with Alpaca. Notably, under Alpaca calibration, FLAT-LLM maintains high accuracy, with less than a 5% drop observed at up to 30% compression. These results highlight the robustness of FLAT-LLM across diverse tasks, calibration datasets, and compression regimes.

6 Conclusion

We propose a fast and accurate training-free structural compression method for LLMs, leveraging low-rank transformations in the activation space. By combining an importance-preserving global rank allocation strategy with efficient head-wise PCA-based approximations, our approach delivers model with strong generation performance using minimal calibration time. Experiments on LLaMA and Mistral show superior generalization and substantial inference speedups compared to prior structural pruning and decomposition-based methods.

Limitations

While FLAT-LLM demonstrates strong empirical performance in compressing large language models without fine-tuning, several limitations remain.

First, although FLAT-LLM achieves superior performance and reduced memory consumption when combined with post-training quantization methods such as GPTQ, further CUDA kernel-level optimization is required to fully realize the potential inference speedup in practice. Second, while FLAT-LLM attains substantial throughput acceleration even at a 10% compression ratio due to its non-uniform rank selection strategy, we do not explicitly analyze how the resulting rank distribution influences end-to-end speedup, nor do we design a hardware-aware rank allocation scheme. We leave both directions for future work.

Acknowledgments

This work is co-funded by Intel Strategic Research Sectors (SRS) - Systems Integration SRS & Devices SRS.

References

- Joshua Ainslie, James Lee-Thorp, Michiel de Jong, Yury Zemlyanskiy, Federico Lebrun, and Sumit Sanghai. 2023. Gqa: Training generalized multi-query transformer models from multi-head checkpoints. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 4895–4901.
- Yongqi An, Xu Zhao, Tao Yu, Ming Tang, and Jinqiao Wang. 2024. Fluctuation-based adaptive structured pruning for large language models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 10865–10873.
- Saleh Ashkboos, Maximilian L Croci, Marcelo Gennari do Nascimento, Torsten Hoefler, and James Hensman. 2024. Slicegpt: Compress large language models by deleting rows and columns. In *The Twelfth International Conference on Learning Representations*.
- Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, and 1 others. 2023. Qwen technical report. *arXiv preprint arXiv:2309.16609*.
- Leo Gao, Jonathan Tow, Baber Abbasi, Stella Biderman, Sid Black, Anthony DiPofi, Charles Foster, Laurence Golding, Jeffrey Hsu, Alain Le Noac’h, Haonan Li, Kyle McDonell, Niklas Muennighoff, Chris Ociepa, Jason Phang, Laria Reynolds, Hailey Schoelkopf, Aviya Skowron, Lintang Sutawika, and 5 others. 2024a. [A framework for few-shot language model evaluation](#).
- Shangqian Gao, Ting Hua, Yen-Chang Hsu, Yilin Shen, and Hongxia Jin. 2024b. Adaptive rank selections for low-rank approximation of language models. In *Proceedings of the 2024 Conference of the North*

- American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 227–241.
- Wei Huang, Yangdong Liu, Haotong Qin, Ying Li, Shiming Zhang, Xianglong Liu, Michele Magno, and Xiaojuan Qi. 2024. Billm: pushing the limit of post-training quantization for llms. In *Proceedings of the 41st International Conference on Machine Learning*, pages 20023–20042.
- Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, L  lio Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timoth  e Lacroix, and William El Sayed. 2023. *Mistral 7b*. Preprint, arXiv:2310.06825.
- Xiaoqi Jiao, Yichun Yin, Lifeng Shang, Xin Jiang, Xiao Chen, Linlin Li, Fang Wang, and Qun Liu. 2020. Tinybert: Distilling bert for natural language understanding. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 4163–4174.
- Chi-Heng Lin, Shangqian Gao, James Seale Smith, Abhishek Patel, Shikhar Tuli, Yilin Shen, Hongxia Jin, and Yen-Chang Hsu. 2024. Modegpt: Modular decomposition for large language model compression. *arXiv preprint arXiv:2408.09632*.
- Aixin Liu, Bei Feng, Bing Xue, Bingxuan Wang, Bochao Wu, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan, and 1 others. 2024. Deepseek-v3 technical report. *arXiv preprint arXiv:2412.19437*.
- Xinyin Ma, Gongfan Fang, and Xinchao Wang. 2023. Llm-pruner: On the structural pruning of large language models. *Advances in neural information processing systems*, 36:21702–21720.
- Xin Men, Mingyu Xu, Qingyu Zhang, Bingning Wang, Hongyu Lin, Yaojie Lu, Xianpei Han, and Weipeng Chen. 2024. Shortgpt: Layers in large language models are more redundant than you expect. *arXiv preprint arXiv:2403.03853*.
- Mingjie Sun, Zhuang Liu, Anna Bair, and J Zico Kolter. 2024. A simple and effective pruning approach for large language models. In *The Twelfth International Conference on Learning Representations*.
- Zhiqing Sun, Hongkun Yu, Xiaodan Song, Renjie Liu, Yiming Yang, and Denny Zhou. 2020. Mobilebert: a compact task-agnostic bert for resource-limited devices. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2158–2170.
- Jiayi Tian, Chao Fang, Haonan Wang, and Zhongfeng Wang. 2023. Bebert: Efficient and robust binary ensemble bert. In *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1–5. IEEE.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timoth  e Lacroix, Baptiste Rozi  re, Naman Goyal, Eric Hambro, Faisal Azhar, and 1 others. 2023. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*.
- Tycho FA van der Ouderaa, Markus Nagel, Mart Van Baalen, and Tijmen Blankevoort. 2024. The llm surgeon. In *The Twelfth International Conference on Learning Representations*.
- Xin Wang, Yu Zheng, Zhongwei Wan, and Mi Zhang. 2024. Svd-llm: Truncation-aware singular value decomposition for large language model compression. *arXiv preprint arXiv:2403.07378*.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pi  ric Cistac, Tim Rault, R  mi Louf, Morgan Funtowicz, and 1 others. 2019. Huggingface’s transformers: State-of-the-art natural language processing. *arXiv preprint arXiv:1910.03771*.
- Yifan Yang, Kai Zhen, Bhavana Ganesh, Aram Galstyan, Goeric Huybrechts, Markus M  ller, Jonas M K  bler, Rupak Vignesh Swaminathan, Athanasios Mouchtaris, Sravan Babu Bodapati, and 1 others. 2025. Wanda++: Pruning large language models via regional gradients. *arXiv preprint arXiv:2503.04992*.
- Yifan Yang, Jiajun Zhou, Ngai Wong, and Zheng Zhang. 2024a. Loretta: Low-rank economic tensor-train adaptation for ultra-low-parameter fine-tuning of large language models. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 3161–3176.
- Yifei Yang, Zouying Cao, and Hai Zhao. 2024b. Laco: Large language model pruning via layer collapse. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 6401–6417.
- Hao Yu and Jianxin Wu. 2023. Compressing transformers: features are low-rank, but weights are not! In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pages 11007–11015.
- Zhihang Yuan, Yuzhang Shang, Yue Song, Qiang Wu, Yan Yan, and Guangyu Sun. 2023. Asvd: Activation-aware singular value decomposition for compressing large language models. *arXiv preprint arXiv:2312.05821*.
- Jialin Zhao, Yingtao Zhang, and Carlo Vittorio Cannistraci. 2025. Pivoting factorization: A compact meta low-rank representation of sparsity for efficient inference in large language models. In *Forty-second International Conference on Machine Learning*.
- Longguang Zhong, Fanqi Wan, Ruijun Chen, Xiaojun Quan, and Liangzhi Li. 2024. Blockpruner: Fine-grained pruning for large language models. *arXiv preprint arXiv:2406.10594*.

Table 3: Comparison of zero-shot downstream performance with prior importance-based compression methods for Llama-2 7B at 20% compression ratio and Llama-2 13B at 30% compression ratio.

Model	Method	Ratio	WinoG.	HellaS.	ARC-e	ARC-c	PIQA	Avg.	delta
LLaMA-2 7B	Original	0%	69.06	75.99	74.58	46.25	77.91	68.76	0
	LaCo (Yang et al., 2024b)	20%	60.46	54.08	55.39	35.84	68.34	54.82	-13.94
	ShortGPT (Men et al., 2024)		65.90	62.63	56.06	36.09	70.24	58.18	-10.58
	BlockPruner (Zhong et al., 2024)		62.43	65.87	61.07	37.29	74.21	60.17	-8.59
	FLAT-LLM (ours)	20%	67.88	69.24	70.45	41.38	75.35	64.86	-3.90
LLaMA-2 13B	Original	0%	72.22	79.39	79.42	49.06	80.52	72.12	0
	LaCo (Yang et al., 2024b)	25%	59.27	60.44	54.34	34.56	72.42	55.44	-16.68
	ShortGPT (Men et al., 2024)		70.80	67.80	60.35	41.30	72.74	62.60	-9.52
	BlockPruner (Zhong et al., 2024)		66.30	72.20	65.82	41.38	76.93	64.53	-7.59
	FLAT-LLM (ours)	25%	71.35	72.93	73.48	43.94	76.44	67.63	-4.49

Corollary 4.2 [Reconstruction Error of Multi-head Output PCA Projection] *Let $\mathbf{Y}_v = \text{concat}(\mathbf{Y}_v^1, \dots, \mathbf{Y}_v^H)$ be the concatenated output hidden states. The squared Frobenius norm of the reconstruction error satisfies:*

$$\|\mathbf{Y}_v - \tilde{\mathbf{Y}}_v\|_F^2 = \sum_{h=1}^H \sum_{i=r+1}^{d_h} \lambda_i^h,$$

Proof.

$$\begin{aligned} \|\mathbf{Y}_v - \tilde{\mathbf{Y}}_v\|_F^2 &= \|\text{concat}(\mathbf{Y}_v^h - \tilde{\mathbf{Y}}_v^h)\|_F^2 \\ &= \sum_{h=1}^H \|(\mathbf{Y}_v^h - \tilde{\mathbf{Y}}_v^h)\|_F^2 = \sum_{h=1}^H \sum_{i=r+1}^{d_h} \lambda_i^h. \end{aligned}$$

□

Therefore, the reconstruction loss of multi-head value output equals to the sum of the dropped eigenvalues of all heads. Truncating the smallest $d_h - r$ eigenvalues of each head leads to the lowest reconstruction loss.

B.2 Empirical Results

Table 4: Comparison of layer-wise reconstruction error for SVD-LLM and FLAT-LLM under 20% sparsity.

Method	Layer 0	Layer 10	Layer 20
SVD-LLM (Wang et al., 2024)	1.56	1.56	1.53
FLAT-LLM (ours)	1.53	1.52	1.52

To evaluate the reconstruction quality of the attention block, we conducted an empirical study measuring the relative Frobenius norm between the compressed and original attention outputs. For a fair comparison, we used the same value-layer inputs from the original model at each layer for both SVD-LLM and FLAT-LLM under a uniform rank setting, thereby avoiding the influence of inter-layer error propagation. We evaluated the reconstruction

error at a 20% compression ratio for both methods on LLaMA-2 7B. The results show that FLAT-LLM consistently yields lower reconstruction error, indicating that our head-wise PCA design for the attention block more effectively preserves the low-rank structure of the activation space than SVD-LLM.

C Nyström-based Low-rank MLP Approximation

To compress the MLP layers, we apply a structured Nyström approximation guided by data-dependent ridge leverage scores inspired by MoDeGPT (Lin et al., 2024). We further improve its efficiency by eliminating the need for the time-consuming pseudo-inverse computation, reducing the overall runtime from hours in MoDeGPT to only a few minutes in our approach. Let d_{hid} and d_{int} denote the hidden and intermediate dimensions, respectively. The method operates on the up-projection matrix $\mathbf{W}_1 \in \mathbb{R}^{d_{hid} \times d_{int}}$ and the down-projection matrix $\mathbf{W}_2 \in \mathbb{R}^{d_{int} \times d_{hid}}$ within each MLP block. To capture activation structure, we compute the correlation matrix C_σ over intermediate hidden states passed through the SiLU activation. Ridge leverage scores are then derived from the regularized correlation matrix via $C_\sigma(C_\sigma + \mathbf{I})^{-1}$, quantifying the relative importance of each intermediate channel. Based on these scores, we construct a selection matrix $\mathbf{S}_k \in \mathbb{R}^{d_{int} \times k}$ that retains the top- k most informative channels, where k is determined by the target sparsity ratio. The up-projection is compressed by selecting the top- k columns, i.e., $\mathbf{W}_1 \mathbf{S}_k$, while the down-projection is approximated via Nyström reconstruction: $(\mathbf{S}_k^\top C_\sigma \mathbf{S}_k)^{-1} \mathbf{S}_k^\top C_\sigma \mathbf{W}_2$. This data-aware procedure preserves key activation subspaces while significantly reducing parameter count and computation.

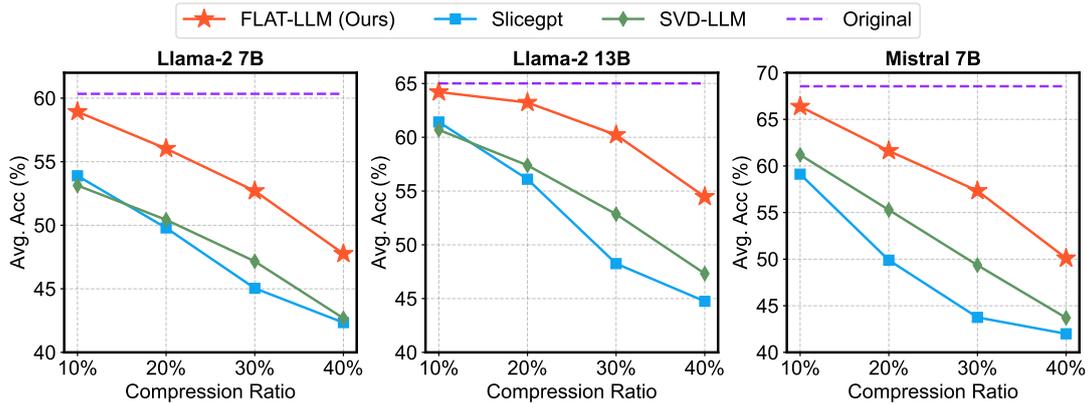


Figure 8: Comparison of zero-shot average accuracy versus compression ratio on various models.

D Additional Experimental Results

D.1 Calibration Efficiency

Table 5: Comparison of calibration time and performance with other methods on LLAMA-2 13B at 20% compression ratio.

Method	Time ↓	PPL ↓	Avg. ACC ↑
SliceGPT (Ashkboos et al., 2024)	0h35m	7.10	0.51
SVD-LLM (Wang et al., 2024)	0h27m	7.69	0.56
FLAP (An et al., 2024)	0h10m	5.90	0.57
FLAT-LLM (ours)	0h15m	5.55	0.63

Table 5 compares the calibration time, perplexity (PPL), and average zero-shot accuracy of various compression methods at a 20% compression ratio on Llama-2 13B. The results are collected from a single A100 GPU. Among the evaluated methods, FLAT-LLM achieves the best overall performance, attaining the lowest perplexity and the highest average accuracy, while maintaining a moderate calibration time of 15 minutes. In contrast, SliceGPT and SVD-LLM exhibit 1.8 – 2.3× longer calibration times and significantly 7 – 12% accuracy drop. Although FLAP achieves the shortest calibration time, it suffers from a 6% accuracy gap compared to our FLAT-LLM. These results highlight that FLAT-LLM offers the best trade-off between calibration efficiency and compression performance, demonstrating high practical deployability.

D.2 Additional Evaluation on Language Modeling

We further compare our FLAT-LLM against recent structural pruning methods, including Pivoting Factorization (Zhao et al., 2025) and LLM Surgeon (van der Ouderaa et al., 2024). All methods are calibrated using 128 samples from WikiText-2, and perplexity is evaluated on the WikiText-2 test split.

Table 6: Performance on LLaMA-3 8B evaluated on WikiText-2.

Method	20%	30%	40%
Pivoting Factorization (Zhao et al., 2025)	8.31	10.83	16.41
LLM Surgeon (van der Ouderaa et al., 2024)	11.27	18.80	52.97
FLAT-LLM (ours)	8.15	9.52	11.67

As shown in Table 6, FLAT-LLM consistently outperforms both Pivoting Factorization and LLM Surgeon on LLaMA-3 8B across various sparsity levels. In addition to strong performance, FLAT-LLM is significantly more efficient: it requires only 8 minutes of calibration on a single A100 40GB GPU, whereas LLM Surgeon takes approximately 4 hours using 8×A100 40GB GPUs when $n_{shots} = 4$.

D.3 Additional Evaluation on Zero-shot Downstream tasks

Figure 8 presents a comprehensive comparison of average zero-shot accuracy across compression ratios ranging from 10% to 40%, evaluated on three LLM models: LLaMA-2 7B, LLaMA-2 13B, and Mistral-7B. Calibration is performed using 128 samples with a sequence length of 4096 from the Alpaca dataset. The performance of the uncompressed dense model is shown as the upper bound (purple dashed line). Our proposed method, FLAT-LLM, consistently outperforms prior low-rank decomposition approaches, including SlicedGPT and SVD-LLM, across all models and compression levels. Notably, FLAT-LLM maintains high performance with less than a 2% accuracy drop at 10% compression across all models. These results underscore the effectiveness and scalability of FLAT-LLM across varying models and compression regimes.

We further evaluate FLAT-LLM against addi-

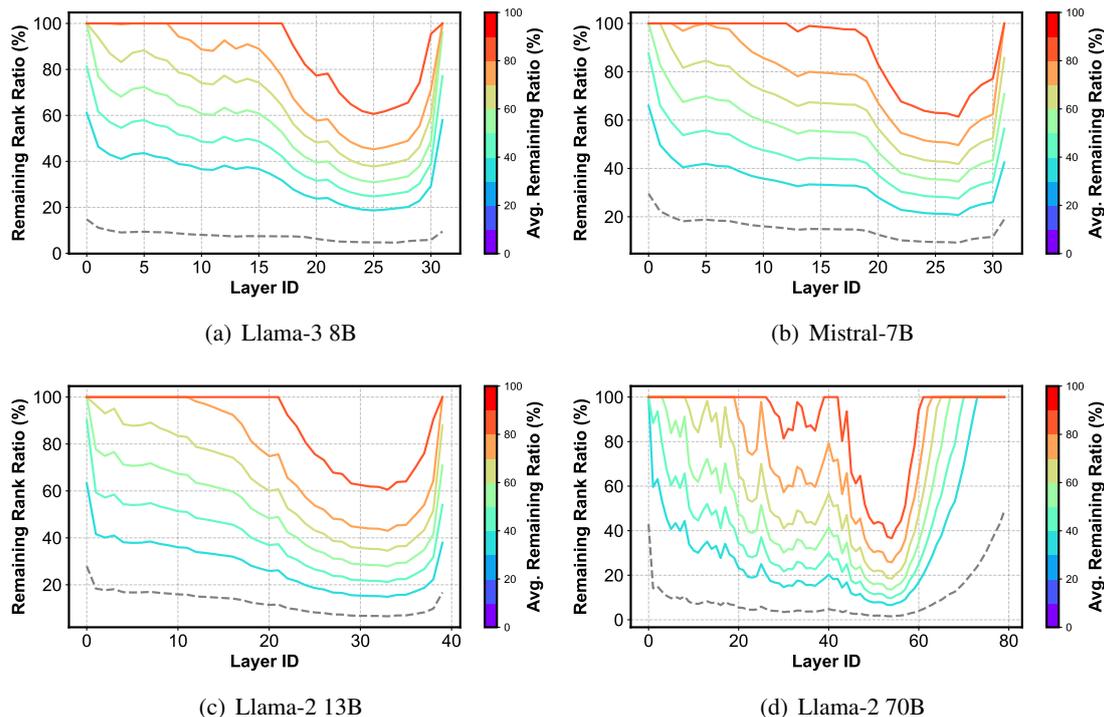


Figure 9: Layer-wise sparsity score visualization across different LLM models using our IPRS algorithm.

tional importance-based baselines on five reasoning tasks, as shown in Table 3. Following the setting of BlockPurmer, we use 256 samples on Alpaca datasets for calibration. On average, FLAT-LLM incurs less than a 4% accuracy drop when pruning Llama-2 7B at a 20% compression ratio, and approximately a 6% drop on Llama-2 13B at a 30% compression ratio—significantly outperforming all baselines. Moreover, FLAT-LLM consistently achieves the highest accuracy across all benchmarks, demonstrating strong generalization across tasks.

D.4 Illustration for IPRS Rank Distribution

Figure 9 presents the layer-wise remaining rank ratios produced by the IPRS algorithm across various LLM models, including Llama-3 8B, Llama-2 13B, 70B, and Mistral-7B. For each model, curves are shown under different target average remaining ratios, ranging from 30% to 90%. To compute the importance of each model, we use 128 samples with 4096 sequence length randomly selected from WikiText-2 dataset. The gray dashed line represents the normalized importance scores t used to guide the adaptive rank allocation. Across all models, IPRS consistently preserves more capacity in layers deemed more important, while aggressively pruning less critical layers. Notably, the resulting

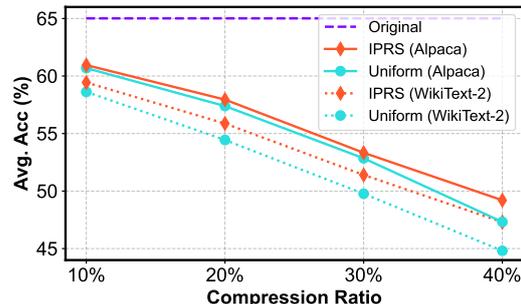


Figure 10: Comparison of average zero-shot accuracy across eight downstream datasets versus compression ratio using uniform rank and our IPRS algorithm on SVD-LLM.

rank distributions exhibit a U-shaped or skewed U-shaped pattern, allocating higher ranks to early, middle, or final layers depending on model-specific importance trends. Despite variations in architecture and depth, all models share a common pattern in which the input-adjacent and output-adjacent layers are more heavily preserved, reflecting their broader importance for information transformation and representation. This consistent behavior across diverse models highlights the robustness and generalization ability of the IPRS algorithm in learning effective, non-uniform rank allocations tailored to model-specific importance profiles.

D.5 IPRS on Other Decomposition Methods

To evaluate the generality of our rank selection method, we also apply IPRS to SVD-LLM, as shown in Figure 10. While it leads to improvements in this setting, the average accuracy gain is modest, typically up to 2.5%. This is due to the distinct truncation loss patterns across layers: SVD-LLM exhibits relatively uniform truncation loss, whereas FLAT-LLM displays highly variable loss, amplifying the benefit of adaptive rank allocation. Overall, combining IPRS with our head-wise PCA-based compression in FLAT-LLM yields consistently superior performance, underscoring the complementary strengths of fine-grained PCA and importance-aware rank selection.

D.6 Comparison with Other Rank-Adaptive Methods

Table 7: Comparison of FLAT-LLM with IPRS and ModeGPT rank allocation strategies on LLaMA-3 8B evaluated on WikiText-2.

Method	30%	40%	50%
FLAT-LLM (w/ ModeGPT’s rank allocation, $\epsilon = 0.2$)	9.66	15.51	81.86
FLAT-LLM (w/ ModeGPT’s rank allocation, $\epsilon = 1$)	62.16	115.86	24.19
FLAT-LLM (w/ our IPRS)	9.52	11.67	16.58

We claim that our IPRS algorithm is entirely hyperparameter-free and achieves higher accuracy than MoDeGPT’s rank allocation method. In Table 7, we apply both IPRS and MoDeGPT’s allocation strategies to FLAT-LLM and evaluate the resulting perplexity on the WikiText-2 test split. The results show that IPRS consistently outperforms MoDeGPT’s allocation method across all compression ratios and hyperparameter settings of ϵ . Furthermore, unlike MoDeGPT, our approach avoids the instability caused by the sensitivity of its hyperparameter ϵ .