Optimizing Singular Spectrum for Large Language Model Compression

Dengjie Li*, Tiancheng Shen*, Yao Zhou, Bernard Ghanem, Yibo Yang, Ming-Hsuan Yang

arXiv:2502.15092v2 [cs.CL] 16 Oct 2026

Abstract

Large language models (LLMs) have demonstrated remarkable capabilities, yet prohibitive parameter complexity often hinders their deployment. Existing singular value decomposition (SVD) based compression methods simply deem singular values as importance scores of decomposed components. However, this importance ordered by singular values does not necessarily correlate with the performance of a downstream task. In this work, we introduce SoCo (**S**ingular spectrum **o**ptimization for large language model **Co**mpression), a novel compression framework that learns to rescale the decomposed components of SVD in a data-driven manner. Concretely, we employ a learnable diagonal matrix to assign importance scores for singular spectrum and develop a three-stage training process that progressively refines these scores—from initial coarse compression to fine-grained sparsification—thereby striking an effective balance between aggressive model compression and performance preservation. Thanks to the learnable singular spectrum, SoCo adaptively prunes components according to the sparsified importance scores, rather than relying on the fixed order of singular values. More importantly, the remaining components with amplified importance scores can compensate for the loss of the pruned ones. Experimental evaluations across multiple LLMs and benchmarks demonstrate that SoCo surpasses the state-of-the-art methods in model compression.

1 Introduction

Recent advances in large language models (LLMs) have led to remarkable breakthroughs in natural language processing (NLP) tasks. (Waswani et al., 2017; Devlin, 2018; Brown et al., 2020; Kaplan et al., 2020; Touvron et al., 2023a) However, the vast number of parameters in these deep transformer architectures results in prohibitive storage and computational requirements and poses significant challenges for deployment and inference efficiency (Wan et al., 2023; Zhu et al., 2023). Consequently, compression of LLMs has emerged as a vital problem (Dettmers et al., 2022; Ma et al., 2023; Zhong et al., 2024b; Hsu et al., 2022).

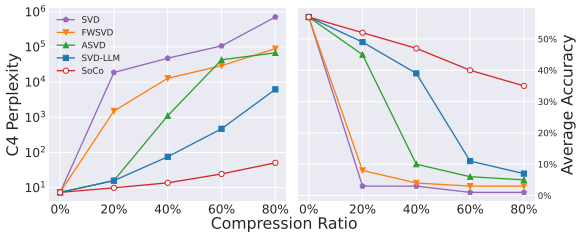

Figure 1: SoCo consistently outperforms existing methods across a range of compression ratios, yielding lower perplexity on C4 (left) and higher average classification accuracy on LM-Evaluation-Harness (right).

A classical approach to model compression performs singular value decomposition (SVD) (Stewart, 1993; Wall et al., 2003) for weight matrices and only preserves the principal components with the largest several singular values, which consist of low-rank matrices and largely reduce the number of parameters. These SVD-based methods assume that the decomposed components with smaller singular values are less important for model performance. However, the order of singular values only reflects their contribution to low-rank approximation but does not necessarily correlate with the performance of a downstream task. To address this, recent works, such as ASVD (Yuan et al., 2023) and SVD-LLM (Wang et al., 2024), introduce activation awareness with calibration data into the decomposition process to mitigate such mismatch. Nonetheless, these methods still perform truncation by the order of singular values without re-evaluating the intrinsic importance of their corresponding components. As a result, the model performance will be rapidly unacceptable when the compression ratio increases. For exam-

*Equal contribution.

ple, as shown in Fig. 1, when compression ratio is larger than 40%, all the compared methods including original SVD (Wall et al., 2003), FWSVD (Hsu et al., 2022), ASVD (Yuan et al., 2023), and SVD-LLM (Wang et al., 2024) lead to catastrophic perplexity on C4 (Raffel et al., 2019). Besides, most existing methods cannot compensate for the loss caused by the pruned components (Hsu et al., 2022; Yuan et al., 2023).

In this paper, we propose **SoCo** (**S**ingular spectrum **o**ptimization for large language model **Co**mpression), a novel framework that learns to rescale the singular spectrum in a data-driven manner. Concretely, SoCo introduces a learnable diagonal matrix in which each element contributes an importance score to re-evaluate the corresponding component. We develop a three-stage training process to deal with the competing objectives of aggressive parameter reduction and performance preservation. In Stage 1, the model undergoes rapid, coarse compression to establish initial importance scores and reach the target compression ratio. Stage 2 uses an alternating optimization strategy to refine these scores, allowing the borderline components to be thoroughly evaluated through controlled oscillation around the target ratio. Stage 3 enforces sparsity in the scores, clearly distinguishing essential components from those that can be pruned.

The whole training process is efficient since we keep all the model parameters frozen. After training, we prune the components according to the sparsified importance scores, rather than simply truncating by the order of singular values as adopted in prior studies. This adaptive, data-aware strategy overcomes the task-agnostic evaluation of singular values and allows for a nuanced selection of the decomposed components. More importantly, the preserved components whose importance scores are amplified can compensate for the loss of pruning, further improving the compression-performance trade-off. The main contributions of this work are:

- We propose SoCo, a learning-based compression framework that optimizes the singular value spectrum through trainable importance scores to overcome the mismatch between the order of singular values from SVD and the true contribution to downstream task performance.
- We design a three-stage training strategy that gradually refines the importance scores through initial coarse compression, midterm adjustment, and final sparsification, resulting in a superior compression-performance trade-off. The preserved components with amplified importance scores can further compensate for the performance loss caused by pruning.
- Extensive evaluations with six popular LLMs (from 7B to 30B parameters) across various benchmarks in Tab. 2 demonstrate that SoCo achieves better compression performance than the state-of-the-art methods. Notably, compared with SVD-LLM (Wang et al., 2024), SoCo gains improvements of up to 37.6%, 81.8%, 94.8%, and 99.2% in performance at 20%, 40%, 60%, 80% compression rates on the C4 dataset (Raffel et al., 2019).

2 Related Work

Large Language Model Compression. LLM compression methods aim to make LLMs more efficient for deployment (Wan et al., 2023; Zhu et al., 2023). Common techniques include quantization, pruning, knowledge distillation, and low-rank decomposition. Quantization reduces the precision of model weights and activations, thereby decreasing memory usage and computational load. (Dettmers et al., 2022; Kim et al., 2024; Shen et al., 2024; Lin et al., 2024). Pruning involves removing less significant weights or neurons from the model to create a sparser architecture. However, pruning often necessitates retraining to recover potential performance degradation (Ma et al., 2023; Ashkboos et al., 2024; Zhong et al., 2024a; Hsieh et al., 2023). Knowledge distillation transfers knowledge from a large "teacher" model to a smaller "student" model by training the latter to replicate the former's outputs. This approach enables the student model to achieve performance comparable to the teacher model while being more efficient (Zhong et al., 2024b; Muralidharan et al., 2024; Hsieh et al., 2023). Low-rank decomposition compresses weight matrices by factorizing them into smaller, computationally efficient components (Hsu et al., 2022; Yuan et al., 2023; Wang et al., 2024).

Low-rank Decomposition Compression. A significant portion of existing low-rank compression techniques (Li et al., 2023; Noach and Goldberg, 2020; Chen et al., 2023; Hsu et al., 2022; Yuan et al., 2023; Wang et al., 2024) relies on Singular Value Decomposition (SVD) (Stewart, 1993; Wall et al., 2003) to factorize and compress model parameters. Although original SVD (Stewart, 1993; Wall et al., 2003) decomposes weight

matrices into lower-rank components for reconstruction, its objective may lead to information loss if not tailored for downstream tasks. Improved SVD-based methods address these limitations. For example, FWSVD (Hsu et al., 2022) leverages Fisher information to preserve critical components, ASVD (Yuan et al., 2023) scales singular values based on activation sensitivity, and SVD-LLM (Wang et al., 2024) incorporates data whitening and layer-wise updates to compensate for accuracy loss at high compression ratios.

However, these approaches depend on a fixed descending order of singular values from SVD, allowing only truncation adjustments without re-evaluating the true significance of each decomposed component. In contrast, our method introduces a learnable mechanism that re-evaluates component importance end-to-end, enabling a more adaptive and effective compression strategy.

3 Method

In Section 3.1, we review LLM architectures and SVD-based compression. Section 3.2 details the SoCo (Singular spectrum optimization for large language model Compression) framework, while Section 3.3 outlines our three-stage training for refining importance scores and rescaling decomposed spectrum. Finally, Section 3.4 examines the learning dynamics during training, demonstrating how SoCo preserves essential model information while effectively reducing parameters.

3.1 Preliminaries

Large language models are built on transformer layers that rely on Multi-Head Self-Attention (MHSA) mechanisms. In each layer, the input sequence is transformed into queries, keys, and values via learned linear transformations to compute attention, followed by processing through a feed-forward network (FFN). The weight matrices W (used for query, key, value, and output projections in attention, as well as up, down, and gate projections in the FFN) comprise the majority of the model's parameters. These high-dimensional matrices often exhibit redundancy, with key information concentrated in a few decomposed components (Haink, 2023; Han et al., 2015), making them the primary targets for compression. This phenomenon can be formally captured using SVD, which factorizes a weight matrix $W \in \mathbb{R}^{m \times n}$ as follows:

$$W = U \Sigma V^{\top}, \tag{1}$$

where $\mathbb{R}$ denotes real numbers, $U \in \mathbb{R}^{m \times min(m,n)}$ and $V \in \mathbb{R}^{min(m,n) \times n}$ are orthogonal matrices and $\Sigma \in \mathbb{R}^{min(m,n) \times min(m,n)}$ is a diagonal matrix containing the singular values in descending order. The magnitude of each singular value quantifies the amount of energy in its corresponding direction, effectively identifying the principal decomposed components of the matrix (Stewart, 1993; Wall et al., 2003). By preserving only the top k largest singular values and their associated vectors, SVD provides a low-rank approximation:

$$W \approx U_k \Sigma_k V_k^{\top}, \tag{2}$$

which preserves the essential features of the parameter matrices. This property is useful for LLM compression, as it reduces LLM's storage and computational requirements when $k \ll \min(m, n)$.

3.2 Overall Framework

Although recent approaches such as ASVD (Yuan et al., 2023) and SVD-LLM (Wang et al., 2024) introduce compression-specific strategies, they still rely on truncating by the descending order of singular values from SVD. Since the original singular values only reflect their contributions to low-rank approximation, this rigid truncation is not necessarily in accordance with the true importance of each decomposed component for downstream tasks and may lead to suboptimal performance at higher compression ratios.

To address these limitations, we propose SoCo (**S**ingular spectrum **o**ptimization for large language model **Co**mpression), a learning-based framework that adaptively re-evaluates and adjusts the singular spectrum. In SoCo, a trainable diagonal matrix $S \in \mathbb{R}^{min(m,n) \times min(m,n)}$ is introduced to assign importance scores to the singular values. Instead of selecting singular values based on their magnitude, the preserved singular values are determined according to the assigned scores in S, as illustrated in Fig. 2. The modified weight matrix W' is defined as

$$W' = U(\Sigma \odot S)V^{\top}, \tag{3}$$

where $\odot$ denotes element-wise multiplication. After training, S dynamically adjusts the importance distribution of the decomposed components based on their true influences on model performance, thereby improving the trade-off between model compression and performance preservation.

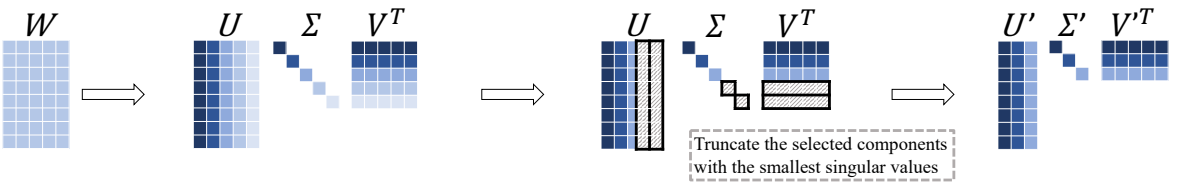

(a) Existing SVD-based compression

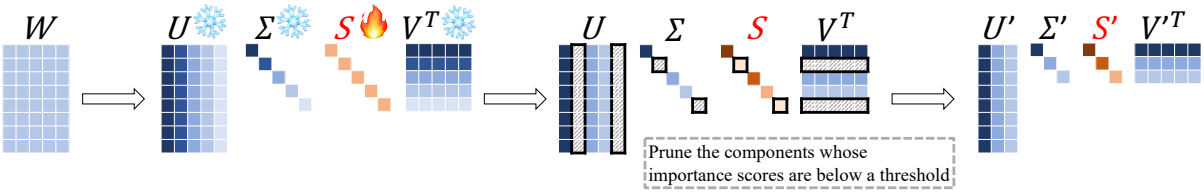

(b) **S**ingular spectrum **o**ptimization for **Co**mpression (SoCo)

Figure 2: Illustration of the overall framework of SoCo. The pre-trained weight matrix W is decomposed into $U\Sigma V^\top$, where Σ is a diagonal matrix with diagonal elements arranged in descending order. a) Existing SVD-based compression methods truncate the smaller singular values and their corresponding vectors in U and rows in $V^\top$. b) The proposed SoCo assigns a diagonal matrix S as importance scores to singular values in Σ. After training, singular values with an importance score below a given threshold (e.g., 0.5) are pruned. In particular, singular values with importance scores larger than the threshold rescale the preserved singular values to compensate the loss from pruned components.

We define the importance scores S as:

$$S = \frac{\lambda_m}{1 + e^{-\lambda_s z + \ln(2\lambda_m - 1)}}, \tag{4}$$

where z is a learnable diagonal matrix (with the same shape as Σ), λ_m controls the overall numerical scale of S (thereby determining the magnitude of re-evaluation), and λ_s modulates the steepness of this sigmoid-like function (larger λ_s results in a more rapid response to changes in z). The term $\ln(2\lambda_m - 1)$ is employed to ensure that the function yields a consistent value at $z = 0$. (More details about S are in A.1 of Appendix.)

To compensate for any deviation introduced by modifying weight matrix W', we further add a trainable deviation term d after the linear transformation of W'. The original model parameters, including the singular values Σ and singular vectors U and V, are frozen during our training.

3.3 Three-stage Optimization

Instead of a single-stage training approach, we propose a three-stage process to adaptively refine the importance scores assigned to the decomposed SVD components. Our approach is motivated by two key considerations. 1) Balancing conflict optimization objectives: compression and performance are inherently conflicting. Excessive parameter reduction can degrade performance while maintaining accuracy often requires retaining more parameters. Single-stage training must simultaneously optimize them, which is intractable to attain a satisfactory balance and may lead to suboptimal minima. 2) Empirical observation of importance score distributions: As shown in Fig. 3, the importance scores after Stage 1 (directly optimizing both compression and performance simultaneously) exhibit a relatively uniform and continuous distribution. As a result, pruning with a fixed threshold may remove many singular values near the threshold that still exert significant influence; when these values are pruned simultaneously across multiple layers, their cumulative effect can substantially degrade performance. Therefore, a more thorough evaluation of the importance of each singular value is needed, along with polarized and sparsified importance scores to prune components with minimal performance degradation.

Figure 3: Distribution of importance scores after each of our three training stages, illustrating the dynamic re-evaluation process.

In light of these considerations, our three-stage training process first rapidly achieves the compres-

Table 1: Conditions and loss functions for each training stage. $step$ denotes the procedure of the whole training process, where $step \in [0, 1]$ (with 0 representing the start and 1 representing the end).

Stage	Conditions	Loss Functions
1	$step < 0.5, r \geq R,$ $until\ r < R$	$L_{inc} + L_{dec}$
2	$step < 0.5, r < R$	L_{inc}
	$step < 0.5, r \geq R$	$L_{inc} + L_{dec}$
3	$step \geq 0.5$	$L_{inc} + L_{spa}$

sion ratio target in stage 1, then polarizes the importance score distribution to gradually separate essential components from redundant ones through alternating optimization in stage 2, and finally sparsifies the importance scores in stage 3. As shown in Fig. 3, after our three-stage training, there is a flat basin between the polarized scores, which minimizes the influence of using a proper threshold to prune the small scores. The process is divided into three stages, each governed by a composite loss function. The corresponding conditions and loss functions are summarized in Tab. 1 and will be detailed in the following subsections.

3.3.1 Stage 1: Rapid and Rough Compression

At the beginning of Stage 1, the current compression ratio $r = 1.0$ exceeds the target compression ratio $R < 1.0$. In this stage, both L_{dec} and L_{inc} jointly drive r from 1.0 to R. The compression ratio r, which is also used as L_{dec}, is defined as:

$$L_{dec} = r = \frac{\sum_{l=1}^{L} (in_l + out_l) \times c_l}{\sum_{l=1}^{L} (in_l \times out_l)}, \tag{5}$$

where in_l and out_l denote the input and output dimensions of layer l, L denotes the number of transformer layers in certain LLM, and c_l is the number of selected decomposed components to preserve in the layer. When working as L_{dec}, to enable gradient flow through the hard thresholding operation, we use the Straight-Through Estimator (STE) (Bengio et al., 2013; Courbariaux et al., 2015):

$$\begin{aligned} c_l &= \sum_{d=1}^{D} \mathbf{1}\left(S_l^d \geq 0.5\right) \\ &= \sum_{d=1}^{D} \mathbf{1}\left(S_l^d \geq 0.5\right) - sg[S_l^d] + S_l^d, \end{aligned} \tag{6}$$

where S_l^d is the importance score of the d-th singular value of layer l and $sg[\cdot]$ halts gradient propagation. $\mathbf{1}(\cdot)$ represents the indicator function, which returns 1 if the condition is true and 0 otherwise. The loss L_{dec} thus applies a positive gradient to reduce S, favoring compression.

Simultaneously, L_{inc} counteracts excessive compression to maintain performance. We define

$$\mathcal{L}_{inc} = \frac{1}{T} \sum_{t=1}^{T} \left(\sum_{c=1}^{C} y_t^c \log \frac{y_t^c}{p_t^c} \right), \tag{7}$$

where T is the number of tokens at the current training step, C is the vocabulary size, y_t^c and p_t^c are the predicted probability distributions before and after model compression. This KL-divergence loss (Cobbe et al., 2021) forces the model to preserve performance during compressing. In practice, the gradient of L_{inc} is smaller than L_{dec}'s gradient in Stage 1. Consequently, L_{dec} becomes the dominant factor, driving rapid model compression until $r < R$. At that point, Stage 2 begins.

3.3.2 Stage 2: Alternating Optimization

After Stage 1, the distribution of importance scores in Fig. 3 shows no clear cutoff for selecting critical components, rendering fixed-threshold pruning inefficient. In Stage 2, we employ an alternating optimization strategy: when $r < R$, L_{dec} is inactive so that L_{inc} enables the recovery of critical components that are pruned too aggressively; when $r \geq R$, L_{dec} is reactivated. This alternation induces controlled oscillations around the target ratio R. With continued training and increased data exposure, the importance assessments gradually stabilize. Such re-evaluation leading to a better singular spectrum for compression (see Tab. 5).

3.3.3 Stage 3: Importance Sparsity

Following rapid convergence in Stage 1 and alternating optimization in Stage 2, enforcing sparsity in the importance scores becomes essential for clearly distinguishing essential components from redundant ones. In Stage 3, we introduce a sparsity loss L_{spa} to drive $S_l^d \leq 0.5$ toward 0 (to prune) and $0.5 < S_l^d \leq 1.0$ toward 1 (to preserve):

$$\mathcal{L}_{spa} = \frac{1}{LD} \sum_{l=1}^{L} \sum_{d=1}^{D} spa_l^d, \tag{8}$$

where the quadratic penalty spa_l^d is defined as

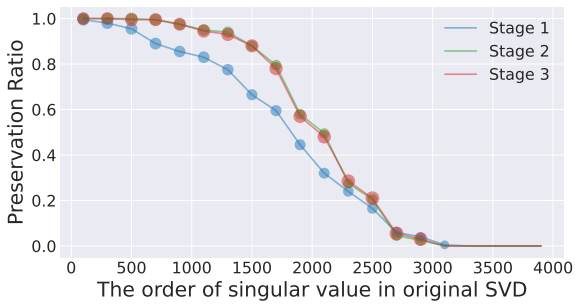

Figure 4: Preservation ratio of components in the original SVD-based importance order across the three stages.

$$spa_l^d = \begin{cases} (S_l^d)^2, & \text{if } S_l^d \leqslant 0.5, \\ (S_l^d - 1)^2, & \text{if } 0.5 < S_l^d \leqslant 1.0, \\ 0, & \text{if } 1.0 < S_l^d. \end{cases} \quad (9)$$

spa_l^d further drives S toward a bimodal distribution, as shown in Fig. 3. Combined with the continued optimization of L_{inc}, Stage 3 reduces the contributions of pruned components while reinforcing essential components. This process creates the basin region of red curve in Fig. 3 where an proper importance score threshold can be easily set to preserve performance with minimal accuracy degradation.

3.4 Analysis

As a pioneering work that integrates decomposed component re-evaluation into SVD-based compression, SoCo provides interesting insights. We summarize our findings as follows:

Validation of Three-stage Optimization: As shown in Tab. 5, each training stage contributes to progressive performance gains, highlighting the effectiveness of phased optimization. This phenomenon validates our design propels the model beyond local optima to a more optimal state.

Importance Score Distribution Evolution Fig. 3 illustrates the evolution of importance score distributions across training stages. In Stage 1, the scores are broadly and continuously distributed. In Stage 2, the distribution polarizes toward two opposing numerical extremes. By Stage 3, S is driven further toward a bimodal distribution, with scores converging near 0.0 (discardable) and 1.0 (critical). This separation minimizes ambiguity and limits loss accumulation in threshold-based pruning, as fewer important components fall near the cutoff.

Adaptation for Component Preservation: Notably, the emergence of importance scores exceeding 1.0 (the right peak of red curve in Fig. 3) unveils a novel mechanism in our framework: the combination of sparsity and output-alignment loss in Stage 3 amplify preserved singular values to counterbalance the loss from pruned components. Unlike existing SVD-based compression methods, which truncate decomposed components, our SoCo can rescale the singular spectrum broadening the effectiveness and potential of SVD.

As shown in Fig. 4, the evolution of the top singular value preservation ratios across the three stages reveals the intrinsic adjustments in component importance. This dynamic re-assessment ensures the real importance distribution of decomposed components for model compression.

4 Experiments

4.1 Experiment Settings

Datasets: We use WikiText-2 (Merity et al., 2016) training set as the primary training corpus for our pipeline. For evaluation, we employ the WikiText-2 (Merity et al., 2016) and C4 (Raffel et al., 2019), and LM-Evaluation-Harness (Gao et al., 2024) testing sets to ensure a comprehensive assessment.

Baselines: Our experiments assess seven LLMs spanning five architectural families and three parameter scales: LLaMA-7B/13B/30B (Touvron et al., 2023a), LLaMA2-7B (Touvron et al., 2023b), OPT-6.7B (Zhang et al., 2022), Vicuna-7B (Chiang et al., 2023), and Mistral-7B (Jiang et al., 2023).

Evaluation Metrics:

- **Perplexity (PPL):** on WikiText-2 (Merity et al., 2016) and C4 (Raffel et al., 2019).
- **Task-Specific Accuracy:** Evaluated using LM-Evaluation-Harness (Gao et al., 2024) on:
  - 6 classification tasks: OpenbookQA (Mihaylov et al., 2018), WinoGrande (Sakaguchi et al., 2019), HellaSwag (Zellers et al., 2019), ARC-e (Clark et al., 2018), PIQA (Bisk et al., 2019), and MathQA (Amini et al., 2019). "Average↑" in the tables below denotes the mean accuracy across the six classification tasks.
  - 1 generative task: TruthfulQA (factual consistency) (Lin et al., 2021).

4.2 Main Results

For a comprehensive comparison, we evaluate SoCo against two categories of model compression techniques: SVD-based methods and pruning methods. For fair evaluations, all experiments are conducted without fine-tuning after compression.

Table 2: Comparison of SoCo with other SVD-based compression methods on the LLaMA-7B model across various compression ratios. The training dataset is the WikiText-2 training set, and evaluation is conducted on the WikiText-2, C4, and LM-Evaluation-Harness testing sets.

RATIO	METHOD	WikiText-2↓	C4↓	Openb.↑	ARC_e↑	WinoG.↑	HellaS.↑	PIQA↑	MathQA↑	Average↑	TruQA.↑
0%	Original	5.68	7.34	34%	75%	70%	57%	79%	27%	57%	30%
20%	SVD	20061	18800	5%	4%	1%	3%	2%	3%	3%	0%
	FWSVD	1727	1511	9%	11%	5%	8%	10%	5%	8%	0%
	ASVD	11.14	15.93	29%	53%	64%	41%	68%	17%	45%	21%
	SVD-LLM	7.94	15.84	**31%**	62%	61%	45%	71%	21%	49%	26%
	SoCo	**6.67**	**9.89**	28%	**70%**	**65%**	**51%**	**75%**	**25%**	**52%**	**33%**
40%	SVD	52489	47774	4%	4%	5%	1%	3%	2%	3%	0%
	FWSVD	18156	12847	6%	5%	2%	0%	5%	3%	4%	0%
	ASVD	1407	1109	8%	11%	9%	8%	13%	8%	10%	1%
	SVD-LLM	13.73	75.42	**25%**	33%	**61%**	40%	63%	12%	39%	17%
	SoCo	**8.60**	**13.71**	23%	**63%**	60%	**45%**	**70%**	**23%**	**47%**	**36%**
60%	SVD	105474	106976	1%	3%	1%	0%	1%	2%	1%	0%
	FWSVD	32194	29292	6%	2%	1%	1%	2%	3%	3%	0%
	ASVD	57057	43036	5%	4%	6%	9%	8%	5%	6%	0%
	SVD-LLM	66.62	471.83	10%	5%	17%	10%	21%	4%	11%	1%
	SoCo	**12.20**	**24.53**	**19%**	**48%**	**55%**	**35%**	**63%**	**21%**	**40%**	**39%**
80%	SVD	687291	708243	0%	4%	2%	1%	1%	0%	1%	0%
	FWSVD	96872	89243	1%	2%	0%	6%	9%	0%	3%	0%
	ASVD	80425	67927	4%	3%	3%	7%	10%	1%	5%	0%
	SVD-LLM	1349	6224	7%	1%	12%	10%	7%	6%	7%	0%
	SoCo	**21.03**	**50.86**	**14%**	**36%**	**53%**	**28%**	**57%**	**20%**	**35%**	**0%**

Table 3: Comparison of SoCo with other SVD-based compression methods on various model families at 20% compression ratio. The training dataset is WikiText-2, while evaluation is conducted on the WikiText-2 and LM-Evaluation-Harness's 6 classification task testing sets.

	OPT-6.7B		LLaMA2-7B		Mistral-7B		Vicuna-7B		LLaMA-13B		LLaMA-30B	
METHOD	WikiText-2↓	Average↑	WikiText-2↓	Average↑	WikiText-2↓	Average↑	WikiText-2↓	Average↑	WikiText-2↓	Average↑	WikiText-2↓	Average↑
Original	10.86	52%	5.47	57%	5.25	61%	6.78	56%	5.09	59%	4.10	61%
SVD	66275	3%	18192	9%	159627	3%	18644	5%	946.31	21%	54.11	33%
FWSVD	14559	6%	2360	12%	6357	8%	2758	9%	15.98	43%	20.54	42%
ASVD	82.00	32%	10.10	36%	13.72	32%	16.23	33%	6.74	54%	22.71	44%
SVD-LLM	16.04	41%	8.50	**53%**	10.21	42%	8.41	51%	6.61	54%	5.63	57%
SoCo	**10.44**	**50%**	**6.69**	52%	**6.69**	**54%**	**7.34**	**53%**	**5.92**	**56%**	**5.17**	**58%**

4.2.1 Comparison with SVD-based Methods

In Tab. 2, SoCo consistently outperforms the competing SVD-based approaches (Hsu et al., 2022; Yuan et al., 2023; Wang et al., 2024) across compression ratios of 20%, 40%, 60%, and 80%. Notably, compared with SVD-LLM (Wang et al., 2024), SoCo achieves relative performance improvements of up to 37.6%, 81.8%, 94.8%, and 99.2% on the C4 dataset (Raffel et al., 2019), corresponding to absolute performance gains ranging from 15.84 to 9.89, 75.42 to 13.71, 471.83 to 24.53, and 6224 to 50.86, respectively. Especially at compression ratios of 60% and 80%, where previous SVD-based methods become impractical, SoCo is the only approach capable of delivering robust performance. Next, Tab. 3 compares our SoCo with other SVD-based compression approaches (Hsu et al., 2022; Yuan et al., 2023; Wang et al., 2024) across different model families and sizes at 20% compression ratio. SoCo consistently achieves superior performance, demonstrating its general applicability across diverse model architectures and sizes.

4.2.2 Comparison with Pruning Methods

We compare SoCo with pruning methods, including LLM-Pruner (Ma et al., 2023), SliceGPT (Ashkboos et al., 2024), and BlockPruner (Zhong et al., 2024a), on the LLaMA-7B (Touvron et al., 2023a) at compression ratios of 26%, 33%, 40%, and 47%. For fair comparison, we follow the baseline methods using the Alpaca (Taori et al., 2023) dataset in training, and performance is evaluated using PPL on the WikiText-2 testing set in Tab. 4. SoCo consistently outperforms these pruning methods.

4.3 Ablation Study

We conduct an ablation study on LLaMA-7B (Touvron et al., 2023a).

Effect of Three-Stage Optimization. Tab. 5 reports the performance on the WikiText-2 (Merity et al., 2016) and C4 (Raffel et al., 2019), and LM-Evaluation-Harness (Gao et al., 2024) testing sets.

Table 4: Comparison with pruning methods on the LLaMA-7B model at compression ratios of 26%, 33%, 40%, and 47%. Following the compared methods, we use the Alpaca dataset for training, and the PPL metric is tested on the WikiText-2 testing set.

METHOD	10GB(26%)	9GB(33%)	8GB(40%)	7GB(47%)
LLM-Pruner	9.88	12.21	18.94	21.68
SliceGPT	8.78	12.73	16.39	27.41
BlockPruner	9.40	12.76	19.78	43.05
SoCo$_{alpaca}$	**8.67**	**10.13**	**12.36**	**16.50**

Table 5: Ablation study of the three-stage training process on the LLaMA-7B (Touvron et al., 2023a) model at compression ratio of 20%. Performance is measured using PPL on WikiText-2 and C4, and the average accuracy on LM-Evaluation-Harness's 6 classification tasks.

RATIO	STAGE	WikiText-2↓	C4↓	Average↑
20%	1	16.21	24.72	41.41%
	1+2	7.13	10.76	50.88%
	1+2+3(SoCo)	**6.67**	**9.89**	**52.48%**
40%	1	109.80	178.16	33.54%
	1+2	27.05	57.14	41.21%
	1+2+3(SoCo)	**8.60**	**13.71**	**47.16%**
60%	1	77.51	146.26	32.32%
	1+2	31.26	66.29	33.06%
	1+2+3(SoCo)	**12.20**	**24.53**	**40.22%**
80%	1	69.86	165.45	32.68%
	1+2	44.37	104.46	33.41%
	1+2+3(SoCo)	**21.03**	**50.86**	**34.63%**

The results demonstrate that our three-stage process progressively refines the importance scores, consistently improving performance. The full three-stage SoCo achieves the best trade-off between compression and performance across all compression ratios.

Training Dataset Analysis. To illustrate the task-aware effect on compression performance, we evaluate our method using four training datasets: the first 20.48 million tokens of the C4 (Raffel et al., 2019) set, the Alpaca (Taori et al., 2023) training dataset, the WikiText-2 (Merity et al., 2016) training set, and the PTB (Marcus et al., 1993) training set. Tab. 6 shows that domain-specific training leads to superior performance on its corresponding test set; for example, when using the C4 set for learning the singular spectrum, the compressed model performs best on the C4 test set. These results indicate that task-aware compression further enhances performance. Our SoCo allows for task-aware singular spectrum rescaling when optimal domain performance is desired. In practice, however, the pruned model may not target a specific domain or may lack dedicated training data. Thus, we use WikiText-2 (Merity et al., 2016) as the default training set in our general-setting experiments (Tab. 2 and Tab. 3).

Table 6: Ablation study on domain-specific training datasets for the LLaMA-7B model at 20% compression ratio. The training datasets include the first 20.48M tokens of the C4 training set, the PTB training set, the Alpaca training set, and the WikiText-2 training set. Evaluation is performed on the C4, PTB, Alpaca, WikiText-2 and LM-Evaluation-Harness test sets.

DATASET	C4↓	PTB↓	Alpaca↓	WikiText-2↓	Average↑
C4	**9.02**	13.70	4.76	7.22	52.25%
PTB	11.93	**12.63**	5.05	8.77	51.64%
Alpaca	9.91	16.17	**4.61**	7.70	52.12%
WikiText-2	9.89	13.61	4.90	**6.67**	**52.48%**

Table 7: Trainable parameters, GPU resources, and training time required for different model sizes.

RESOURCE	LLAMA-7B	LLAMA-13B	LLAMA-30B
Trainable parameters	1.84M (2.83‰)	2.87M (2.26‰)	5.59M (1.74‰)
GPU number	1	1	4
Memory per GPU	23G	44G	36G
Training time	2h10m	3h10m	2h45m

Training Time and GPU Cost. Tab. 7 summarizes the trainable parameter counts, GPU requirements, and training durations across various model sizes. Our results indicate that SoCo is highly resource-efficient, requiring only a small number of trainable parameters and enabling rapid singular spectrum optimization with modest GPU resources.

5 Conclusion

In this paper, we introduce SoCo (**S**ingular spectrum **o**ptimization for large language model **Co**mpression), a novel compression framework that leverages a learnable diagonal matrix to adaptively rescale the singular spectrum of weight matrices. Unlike traditional SVD-based truncation methods, SoCo re-evaluates the intrinsic importance of decomposed components and prunes them accordingly. Our three-stage training process progressively refines the importance scores of the decomposed SVD components, enabling a more fine-grained and effective compression strategy. Extensive experiments on multiple large language models and benchmarks demonstrate that SoCo consistently outperforms existing SVD-based and pruning methods, delivering superior performance even at high compression ratios. These results highlight the potential of our adaptive, data-driven approach for model compression, paving the way for more efficient deployment of large language models in resource-constrained environments.

6 Limitations

Our method currently employs a fixed threshold (e.g., 0.5) for pruning based on the importance scores S. Although this threshold is effective in our experiments, it may not fully capture the nuanced variations in component importance across different model architectures or tasks. Future work could explore adaptive thresholding strategies to further enhance compression performance without compromising general applicability.

References

Aida Amini, Saadia Gabriel, Shanchuan Lin, Rik Koncel-Kedziorski, Yejin Choi, and Hannaneh Hajishirzi. 2019. Mathqa: Towards interpretable math word problem solving with operation-based formalisms. In *NAACL*.

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. *ICLR*.

Yoshua Bengio, Nicholas Léonard, and Aaron Courville. 2013. Estimating or propagating gradients through stochastic neurons for conditional computation. *arXiv preprint arXiv:1308.3432*.

Yonatan Bisk, Rowan Zellers, Ronan Le Bras, Jianfeng Gao, and Yejin Choi. 2019. Piqa: Reasoning about physical commonsense in natural language. In *AAAI*.

Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeff Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Ma teusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. Language models are few-shot learners. *NeurIPS*.

Boyu Chen, Hanxuan Chen, Jiao He, Fengyu Sun, and Shangling Jui. 2023. Ternary singular value decomposition as a better parameterized form in linear mapping. *arXiv preprint arXiv:2308.07641*.

Wei-Lin Chiang, Zhuohan Li, Zi Lin, Ying Sheng, Zhanghao Wu, Hao Zhang, Lianmin Zheng, Siyuan Zhuang, Yonghao Zhuang, Joseph E. Gonzalez, Ion Stoica, and Eric P. Xing. 2023. Vicuna: An open-source chatbot impressing gpt-4 with 90%* chatgpt quality.

Peter Clark, Isaac Cowhey, Oren Etzioni, Tushar Khot, Ashish Sabharwal, Carissa Schoenick, and Oyvind Tafjord. 2018. Think you have solved question answering? try arc, the ai2 reasoning challenge. *arXiv preprint arXiv:1803.05457*.

Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. 2021. Training verifiers to solve math word problems. *NeurIPS*.

Matthieu Courbariaux, Yoshua Bengio, and Jean-Pierre David. 2015. Binaryconnect: Training deep neural networks with binary weights during propagations. *NeurIPS*.

Tim Dettmers, Mike Lewis, Younes Belkada, and Luke Zettlemoyer. 2022. Gpt3. int8 (): 8-bit matrix multiplication for transformers at scale. *NeurIPS*.

Jacob Devlin. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.

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, Eric Tang, Anish Thite, Ben Wang, Kevin Wang, and Andy Zou. 2024. A framework for few-shot language model evaluation.

David Haink. 2023. Hessian eigenvectors and principal component analysis of neural network weight matrices. *arXiv preprint arXiv:2311.00452*.

Song Han, Huizi Mao, and William J Dally. 2015. Deep compression: Compressing deep neural networks with pruning, trained quantization and huffman coding. *arXiv preprint arXiv:1510.00149*.

Cheng-Yu Hsieh, Chun-Liang Li, Chih-Kuan Yeh, Hootan Nakhost, Yasuhisa Fujii, Alexander Ratner, Ranjay Krishna, Chen-Yu Lee, and Tomas Pfister. 2023. Distilling step-by-step! outperforming larger language models with less training data and smaller model sizes. *ACL*.

Yen-Chang Hsu, Ting Hua, Sung-En Chang, Qiang Lou, Yilin Shen, and Hongxia Jin. 2022. Language model compression with weighted low-rank factorization. *ICLR*.

Albert Q Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, et al. 2023. Mistral 7b. *arXiv preprint arXiv:2310.06825*.

Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. 2020. Scaling laws for neural language models. *arXiv preprint arXiv:2001.08361*.

Sehoon Kim, Coleman Richard Charles Hooper, Amir Gholami, Zhen Dong, Xiuyu Li, Sheng Shen, Michael W Mahoney, and Kurt Keutzer. 2024.

Squeezellm: Dense-and-sparse quantization. In *ICML*.

Yixiao Li, Yifan Yu, Qingru Zhang, Chen Liang, Pengcheng He, Weizhu Chen, and Tuo Zhao. 2023. Losparse: Structured compression of large language models based on low-rank and sparse approximation. In *ICML*.

Ji Lin, Jiaming Tang, Haotian Tang, Shang Yang, Wei-Ming Chen, Wei-Chen Wang, Guangxuan Xiao, Xingyu Dang, Chuang Gan, and Song Han. 2024. Awq: Activation-aware weight quantization for on-device llm compression and acceleration. *MLSys*.

Stephanie Lin, Jacob Hilton, and Owain Evans. 2021. Truthfulqa: Measuring how models mimic human falsehoods. *ACL*.

Xinyin Ma, Gongfan Fang, and Xinchao Wang. 2023. Llm-pruner: On the structural pruning of large language models. *NeurIPS*.

Mitch Marcus, Beatrice Santorini, and Mary Ann Marcinkiewicz. 1993. Building a large annotated corpus of english: The penn treebank. *Computational linguistics*.

Stephen Merity, Caiming Xiong, James Bradbury, and Richard Socher. 2016. Pointer sentinel mixture models. *ICLR*.

Todor Mihaylov, Peter Clark, Tushar Khot, and Ashish Sabharwal. 2018. Can a suit of armor conduct electricity? a new dataset for open book question answering. In *EMNLP*.

Saurav Muralidharan, Sharath Turuvekere Sreenivas, Raviraj Bhuminand Joshi, Marcin Chochowski, Mostofa Patwary, Mohammad Shoeybi, Bryan Catanzaro, Jan Kautz, and Pavlo Molchanov. 2024. Compact language models via pruning and knowledge distillation. In *NeurIPS*.

Matan Ben Noach and Yoav Goldberg. 2020. Compressing pre-trained language models by matrix decomposition. In *AACL*.

Colin Raffel, Noam M. Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2019. Exploring the limits of transfer learning with a unified text-to-text transformer. *JMLR*.

Keisuke Sakaguchi, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. 2019. An adversarial winograd schema challenge at scale. In *AAAI*.

Haihao Shen, Naveen Mellempudi, Xin He, Qun Gao, Chang Wang, and Mengni Wang. 2024. Efficient post-training quantization with fp8 formats. *MLSys*.

Gilbert W Stewart. 1993. On the early history of the singular value decomposition. *SIAM review*.

Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy Liang, and Tatsunori B Hashimoto. 2023. Alpaca: A strong, replicable instruction-following model. *Stanford Center for Research on Foundation Models. https://crfm. stanford. edu/2023/03/13/alpaca. html*.

Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. 2023a. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*.

Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. 2023b. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.

Michael E Wall, Andreas Rechtsteiner, and Luis M Rocha. 2003. Singular value decomposition and principal component analysis. In *A practical approach to microarray data analysis*, pages 91–109. Springer.

Zhongwei Wan, Xin Wang, Che Liu, Samiul Alam, Yu Zheng, Jiachen Liu, Zhongnan Qu, Shen Yan, Yi Zhu, Quanlu Zhang, et al. 2023. Efficient large language models: A survey. *arXiv preprint arXiv:2312.03863*.

Xin Wang, Yu Zheng, Zhongwei Wan, and Mi Zhang. 2024. Svd-llm: Truncation-aware singular value decomposition for large language model compression. *ICLR*.

A Waswani, N Shazeer, N Parmar, J Uszkoreit, L Jones, A Gomez, L Kaiser, and I Polosukhin. 2017. Attention is all you need. In *NeurIPS*.

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*.

Rowan Zellers, Ari Holtzman, Yonatan Bisk, Ali Farhadi, and Yejin Choi. 2019. Hellaswag: Can a machine really finish your sentence? In *ACL*.

Susan Zhang, Stephen Roller, Naman Goyal, Mikel Artetxe, Moya Chen, Shuohui Chen, Christopher Dewan, Mona Diab, Xian Li, Xi Victoria Lin, et al. 2022. Opt: Open pre-trained transformer language models. *arXiv preprint arXiv:2205.01068*.

Longguang Zhong, Fanqi Wan, Ruijun Chen, Xiaojun Quan, and Liangzhi Li. 2024a. Blockpruner: Fine-grained pruning for large language models. *arXiv preprint arXiv:2406.10594*.

Qihuang Zhong, Liang Ding, Li Shen, Juhua Liu, Bo Du, and Dacheng Tao. 2024b. Revisiting knowledge distillation for autoregressive language models. *ACL*.

Xunyu Zhu, Jian Li, Yong Liu, Can Ma, and Weiping Wang. 2023. A survey on model compression for large language models. *arXiv preprint arXiv:2308.07633*.

A Appendix

A.1 Functions and Equations

Function graph of importance scores S. As shown in Fig. 5 (with $\lambda_m = 2$ and $\lambda_s = 10$ in SoCo), the function S facilitates performance compensation and rapid convergence by tuning λ_m and λ_s. Specifically, $\lambda_m = 2$ confines the overall numerical range of S to $[0, 2]$, allowing S to compensate for the information loss incurred by pruned components with values exceeding 1.

Meanwhile, λ_s modulates the steepness of the sigmoid-like function S. As shown in Eq. 10, the derivative $\frac{dS}{dz}$ is directly proportional to λ_s, indicating that a larger λ_s yields a faster response to changes in z, thereby accelerating training convergence.

$$\frac{dS}{dz} = \frac{\lambda_m \lambda_s (2\lambda_m - 1) e^{-\lambda_s z}}{[1 + (2\lambda_m - 1) e^{-\lambda_s z}]^2}. \tag{10}$$

The term $\ln(2\lambda_m - 1)$ in Eq. 4 is included to ensure that the function yields the same value as the sigmoid function at $z = 0$. This consistency makes it easy to set a threshold for S.

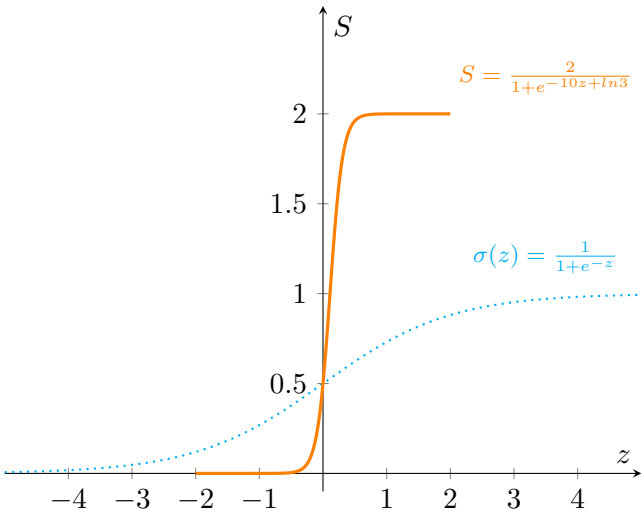

Figure 5: Function graph of the importance scores S. The orange function denotes the S used in SoCo, while the blue function represents the standard sigmoid function.

Function graph of spa in L_{spa}. The function spa in L_{spa} drives each importance score s in S toward 0 when $S \leq 0.5$ and toward 1 when $0.5 < S \leq 1$. Moreover, as illustrated by the orange curve in Fig. 6, spa_l^d imposes a zero gradient for $S > 1$ in Eq. 11, thereby allowing S to exceed 1. This mechanism enables compensation for the information loss induced by pruned components.

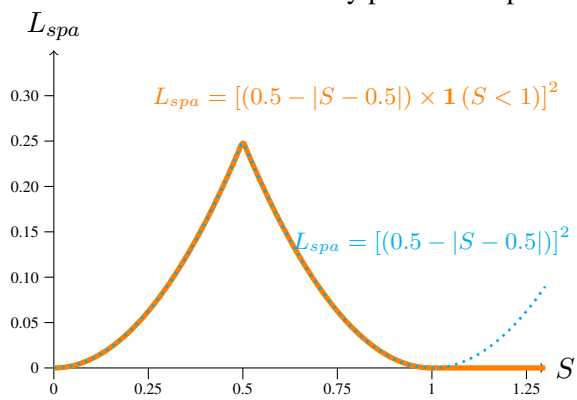

Figure 6: Illustration of the function of spa in $\mathcal{L}_{spa}$ that drive importance scores toward 0 or 1. The orange function used in SoCo converges toward 0 or 1 for scores within the range [0,1] and has no effect on scores above 1. And the blue function ensures that scores greater than 1 are gradually pulled back toward 1.

Therefore, the gradient of L_{spa} with respect to S_l^d is defined piecewise to capture different behaviors depending on the value of S_l^d:

$$\frac{\partial L_{spa}}{\partial S_l^d} = \begin{cases} \frac{2S_l^d}{LD}, & \text{if } S_l^d \leq 0.5, \\ \frac{2(S_l^d - 1)}{LD}, & \text{if } 0.5 < S_l^d \leq 1.0, \\ 0, & \text{if } 1.0 < S_l^d. \end{cases} \tag{11}$$

This formulation encourages sparsity by penalizing values closer to the midpoint of the interval $[0, 1]$, while eliminating gradients for the values larger then 1, thereby promoting the compensation mentioned in the paper.

A.2 Algorithm implementation

Algorithm 1 Removal of singular values whose importance scores $S < 0.5$.

```
#pytorch code
idx = (S>=0.5).nonzero().flatten()
U_k = U.index_select(1, idx)  # U_k stands for Uk in equation (2)
V_k = V.index_select(0, idx)  # V_k stands for Vk in equation (2)
Sigma_k = Sigma.index_select(0, idx)  # Sigma_k stands for Sigmak in equation (2)
S_k = S.index_select(0, idx)  # S_k stands for S', importance score after selection
```

A.3 Implementation Details

Training Settings. The WikiText-2 dataset is concatenated and segmented into samples of 1024 tokens. The hyperparameters λ_m=2.0 and λ_s=10.0 are used as the default values in S. We optimize the models using AdamW with default parameters (weight decay = 0.0, $\beta_1 = 0.9$, $\beta_2 = 0.999$, and $\epsilon = 1 \times 10^{-8}$). Different learning rates for compressing various models are detailed in Tab. 8. A cosine scheduler with a warm-up ratio of 0.05 is applied, and the maximum gradient norm is set to 1.0. Training is performed with a batch size of 1 using BFloat16 precision and gradient checkpointing to optimize GPU memory usage. Most experiments run on a single A100 GPU; however, for LLaMA-30B, we employ FSDP (FULL_SHARD mode) across 4 A100 GPUs.

Learning Rate for Different Settings. To achieve optimal compression results, we select different learning rates based on the model and the target compression ratio. The specific values are provided in Tab. 8.

Table 8: Learning Rate across Different Models and Compression Ratio.

MODEL	Target Compression Ratio	Learning Rate
LLaMA-7B	20%	0.002
LLaMA-7B	40%	0.002
LLaMA-7B	60%	0.007
LLaMA-7B	80%	0.008
LLaMA-13B	20%	0.002
LLaMA-30B	20%	0.003
LLaMA2-7B	20%	0.002
OPT-6.7B	20%	0.002
Mistral-7B	20%	0.001
Vicuna-7B	20%	0.002

A.4 Supplementary Ablation Study

Deviation Term d. To compensate for any deviation introduced by modifying the weight matrix W', we add a trainable deviation term d after the linear transformation of W' in SoCo. This deviation term quantifies the average deviation and helps mitigate performance loss. Its effectiveness is demonstrated in Tab. 9.

Table 9: Ablation experiments of SoCo on the use of the deviation term d, conducted on the LLaMA-7B model at the compression ratios of 20%, 40%, 60%, and 80%.

RATIO	METHOD	WikiText-2↓	C4↓	Average↑
20%	w/o Deviation	6.77	9.92	52.36%
	w Deviation	**6.67**	**9.89**	**52.48%**
40%	w/o Deviation	9.69	15.56	45.52%
	w Deviation	**8.60**	**13.71**	**47.16%**
60%	w/o Deviation	16.27	31.80	38.01%
	w Deviation	**12.20**	**24.53**	**40.22%**
80%	w/o Deviation	33.72	81.57	33.22%
	w Deviation	**21.03**	**50.86**	**34.63%**

Hyper Parameters λ_m and λ_s in S. Excluding one outlier at $(\lambda_m, \lambda_s) = (2.0, 1.0)$ in Tab. 10, the remaining results vary only slightly, demonstrating the robustness of SoCo with respect to these hyperparameters. We adopt a relatively optimal configuration $(\lambda_m, \lambda_s) = (2.0, 10.0)$ for other experiments.

Table 10: Ablation study on the hyperparameters λ_m and λ_s in S. Performance is evaluated via PPL on WikiText-2 and C4, and the averaged accuracy of 6 classification tasks in LM-Evaluation-Harness.

λ_m	λ_s	WikiText-2↓	C4↓	Average↑
1.0	1.0	7.47	12.17	49.38%
1.0	10.0	6.80	9.99	52.42%
2.0	**10.0**	**6.67**	**9.89**	**52.48%**
4.0	10.0	6.75	9.98	51.97%
8.0	10.0	6.77	10.00	51.73%
2.0	1.0	10.30	17.14	42.52%
2.0	2.0	8.46	12.92	49.34%
2.0	5.0	6.98	10.27	51.52%
2.0	15.0	6.73	9.97	52.34%
2.0	20.0	6.77	9.99	52.37%

A.5 Compression Result Observations

Decomposed Components Preservation Ratio Across All Layers in LLaMA at Different Compression Ratios. The four subplots in Fig. 7 show the proportion of preserved components across the layers of LLaMA-7B at compression ratios of 20%, 40%, 60%, and 80%. In each subplot, each point on the horizontal axis represents a transformer layer, and the plot shows seven colored lines—each corresponding to a different linear module's weight matrix within that layer. The vertical values of these lines indicate the preservation proportion of decomposed components from the original SVD for each module.

Overall, the preservation ratio of decomposed components for the FFN linear modules (down_pro, up_proj, gate_proj) is higher than that for the Multi-Head Attention linear modules (q_proj, k_proj, v_proj, o_proj). Moreover, as the compression ratio increases, the decrease in the preservation ratio is less pronounced for the Multi-Head Attention modules than for the FFN modules.

Decomposed Components Preservation Ratio Across All Layers in Different LLMs at 20% Compression Ratio.

Fig. 8 presents the preserved component proportions from SoCo across the transformer layers of various models—LLaMA-7B, LLaMA-30B, Mistral-7B, Vicuna-7B, Llama-2-7b, and OPT-6.7B—all evaluated at a 20% compression ratio. In general, the preservation ratio of decomposed components for the FFN linear modules is higher than that for the Multi-Head Attention linear modules (q_proj, k_proj, v_proj, o_proj). Furthermore, no clear pattern emerges across the various models.

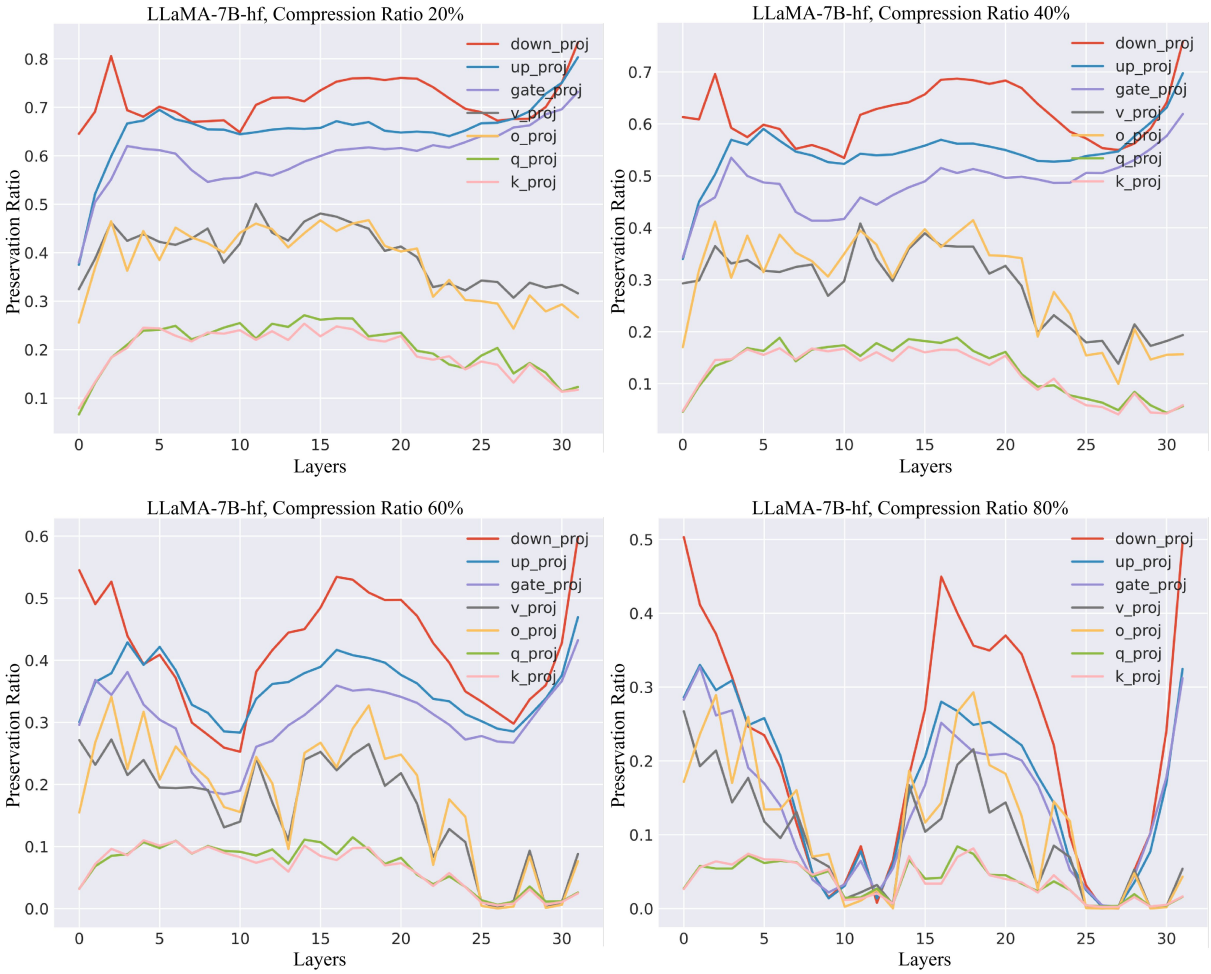

Figure 7: The four subplots depict the preservation ratio of decomposed components in SoCo for LLaMA-7B at the compression ratios of 20%, 40%, 60%, and 80%.

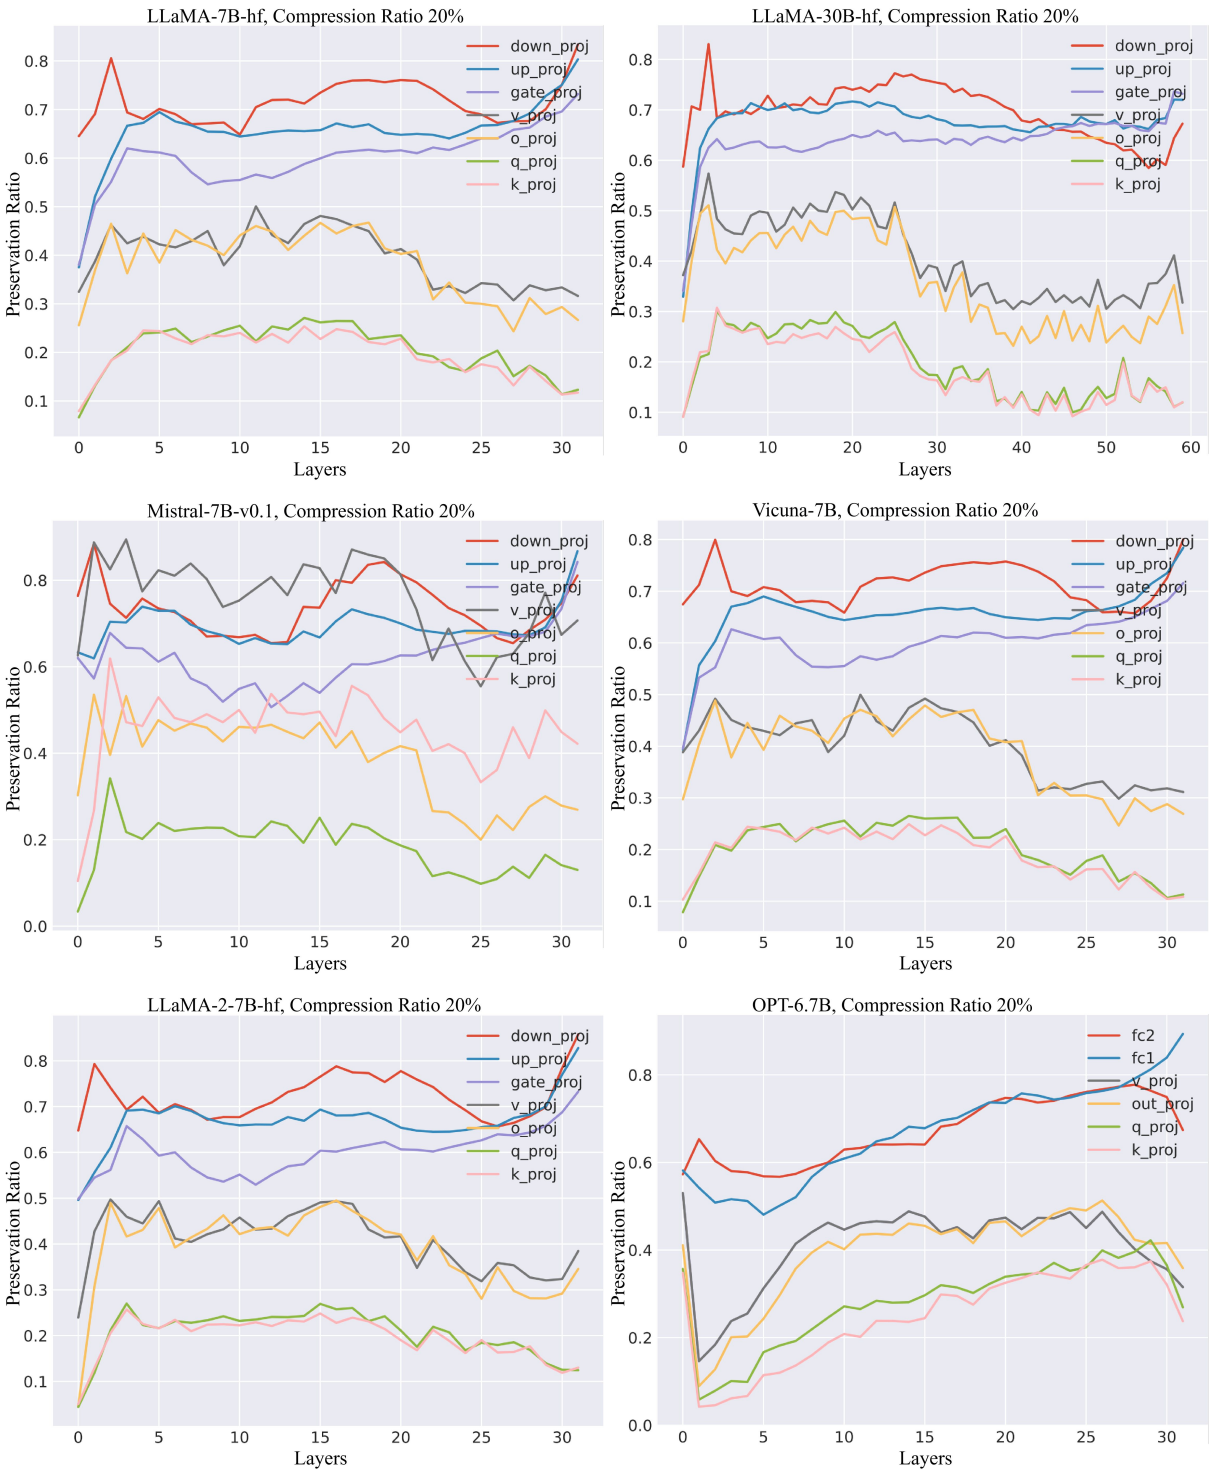

Figure 8: The six subplots depict the preservation ratio of decomposed components in SoCo across all transformer layers and linear modules in various LLMs at 20% compression ratio.