

Feature-based Low-Rank Compression of Large Language Models via Bayesian Optimization

Anonymous ACL submission

Abstract

In recent years, large language models (LLMs) have driven advances in natural language processing. Still, their growing scale has increased the computational burden, necessitating a balance between efficiency and performance. Low-rank compression, a promising technique, reduces non-essential parameters by decomposing weight matrices into products of two low-rank matrices. Yet, its application in LLMs has not been extensively studied. The key to low-rank compression lies in low-rank factorization and low-rank dimensions allocation. To address the challenges of low-rank compression in LLMs, we conduct empirical research on the low-rank characteristics of large models. We propose a low-rank compression method suitable for LLMs. This approach involves precise estimation of feature distributions through pooled covariance matrices and a Bayesian optimization strategy for allocating low-rank dimensions. Experiments on the LLaMA-2 models demonstrate that our method outperforms existing strong structured pruning and low-rank compression techniques in maintaining model performance at the same compression ratio.¹.

1 Introduction

In recent years, the emergence and application of large language models (LLMs) have served as a powerful stimulant for natural language processing and artificial intelligence (OpenAI, 2022, 2023; Bubeck et al., 2023; Yang et al., 2023). Adhering to the scaling law (Kaplan et al., 2020; Hoffmann et al., 2022), researchers are continually seeking LLMs with more parameters and training data, aiming to achieve general models closer to human capabilities. However, larger language models imply a larger overhead of computing resources. Therefore, when deploying LLMs, it is necessary to strike a balance between efficiency and performance (Wan

et al., 2024). To achieve efficient LLMs, many compression techniques for LLMs are proposed, such as pruning (Frantar and Alistarh, 2023a; Sun et al., 2024; Ma et al., 2023), quantization (Frantar et al., 2023; Lin et al., 2023; Liu et al., 2023) and knowledge distillation (Gu et al., 2024).

Among these methods, unstructured pruning and quantization can reduce the number of parameters or memory requirements by half or even more without significant performance degradation, but they require specialized GPU kernels to fully realize their acceleration potential. In contrast, structured pruning can produce lightweight models that do not rely on specialized hardware. Despite extensive research, the performance of structured pruning still lags significantly behind that of the original model. Low-rank compression (LRC) (Ben Noach and Goldberg, 2020; Li et al., 2023) is another promising compression technique. It decomposes the weight matrix into the product of two dense low-rank matrices, discarding unimportant parameter information during the decomposition process. However, LRC remains under-explored in LLMs.

The keys to LRC are low-rank decomposition methods and low-rank dimension allocation. Existing decomposition methods can generally be categorized into two types: weight-based and feature-based decomposition. The former minimizes the reconstruction error of weight matrices by applying truncated SVD or weighted SVD (Ben Noach and Goldberg, 2020; Hsu et al., 2022; Hua et al., 2022). However, recent research (Chen et al., 2021; Yu and Wu, 2023) has discovered that the weights of most Transformer-based language models are typical of high rank or even close to full rank; thus, direct decomposition might result in significant error. In contrast, the model’s features usually exhibit low-rank characteristics. Thus, more work focuses on the feature-based decomposition (Chen et al., 2021; Yu and Wu, 2023; Kaushal et al., 2023), which aims to minimize the reconstruction error

¹All the implementation details and model checkpoints will be publicly available at <https://anonymous.com>.

of features. On the other hand, allocating suitable low-rank dimensions to different weight matrices according to the target compression ratio can also reduce the downside on the model’s overall performance since they exhibit varying sensitivities to low-rank compression.

When LRC is applied to LLMs, it encounters more new challenges. First, it is challenging for LLMs to maintain their generality while achieving feature-based low-rank compression. This is because the feature space of LLMs is extremely high dimensions, making the feature distribution more complex, and the presence of outlier features may interfere with the accurate distribution estimation. Thus, we utilize the pooled covariance matrix instead of the sample covariance matrix, which enables a more accurate estimation of feature distributions (Raninen et al., 2022). Then, for low-rank dimension allocation, manual design struggles to achieve optimal results, and due to its vast search space, grid search requires a considerable amount of time. We conduct empirical studies on the low-rank sensitivity of different types of parameters and observe significant variations among them. Based on these findings, we categorize the parameters into groups, allowing each group to share the same low-rank dimensions. This approach effectively narrows down the search space, and furthermore, we utilize sample-efficient Bayesian optimization to determine the optimal low-rank allocation. To evaluate the effectiveness of our proposed LRC method, we conduct experiments on two commonly used LLaMA-2 models (Touvron et al., 2023). Experimental results demonstrate our proposed method can perform better than existing strong structured pruning and LRC methods in LLMs. When combined with efficient post-training, our method obtains the latest state-of-art for the same settings, maintaining 98% of the model’s performance at the 20% compression rate.

Overall, our main contributions include:

- We analyze the challenges that LLMs face in low-rank compression and, through empirical research, demonstrate that LLMs represented by LLaMA exhibit vastly different sensitivities to low-rank compression across various parameters.
- We propose a novel Bayesian optimization-based feature low-rank compression (Bolaco).
- Extensive experiments show that our Bolaco outperforms the existing strong structured pruning and LRC methods in LLMs.

2 Preliminary

In this section, we summarily introduce the foundation of low-rank factorization in model compression, and then empirically show that different layers of the Transformers-based generative large language model have different low-rank sensitivities.

2.1 Weight-based and Feature-based Low-rank Decomposition

The low-rank decomposition reduces the number of parameters by decomposing the linear layer weights into two low-rank matrices. Weight-based factorization is one naive method. For a linear layer $W \in \mathbb{R}^{d_2 \times d_1}$, the objective can be formulated as:

$$\begin{aligned} \min_{B, A} & \|W - BA\|_F \\ \text{s.t.} & \text{rank}(BA) = r. \end{aligned} \quad (1)$$

If $r < d_1 d_2 / (d_1 + d_2)$, the factorization can reduce the total parameter amount. According to the Eckart–Young–Mirsky theorem, the truncated singular value decomposition (SVD) provides the optimal solution:

$$\begin{aligned} W &= U \Sigma V^T \\ A &= V_r^T, B = U_r \Sigma_r, \end{aligned} \quad (2)$$

where $A \in \mathbb{R}^{r \times d_1}$, $B \in \mathbb{R}^{d_2 \times r}$, Σ_r is the top- r largest singular values, U_r and V_r are the corresponding singular vectors.

However, in the vast majority of cases, the weights of PLMs have a high rank, and a direct truncated SVD decomposition on the weights would lead to significant reconstruction errors (Chen et al., 2021). In comparison, the representation space of PLMs exhibits a clear low-rank property (Yu and Wu, 2023). Therefore, another line of work has considered feature-based factorization:

$$\begin{aligned} \min_{B, A} & \|WX - BAX\|_F \\ \text{s.t.} & \text{rank}(BA) = r. \end{aligned} \quad (3)$$

For the linear layer $Y = WX$, Chen et al. (2021) obtain the optimal solution to Eq. 3 by simultaneously performing the SVD decomposition of the weight and features. Yu and Wu (2023) propose a more efficient Atomic Feature Mimicking (AFM) method, which utilizes the PCA decomposition to find the projection matrices:

$$\begin{aligned} \text{Cov}(Y) &= U \Sigma U^T \\ Y - E[Y] &= U_r U_r^T (WX - E[Y]), \end{aligned} \quad (4)$$

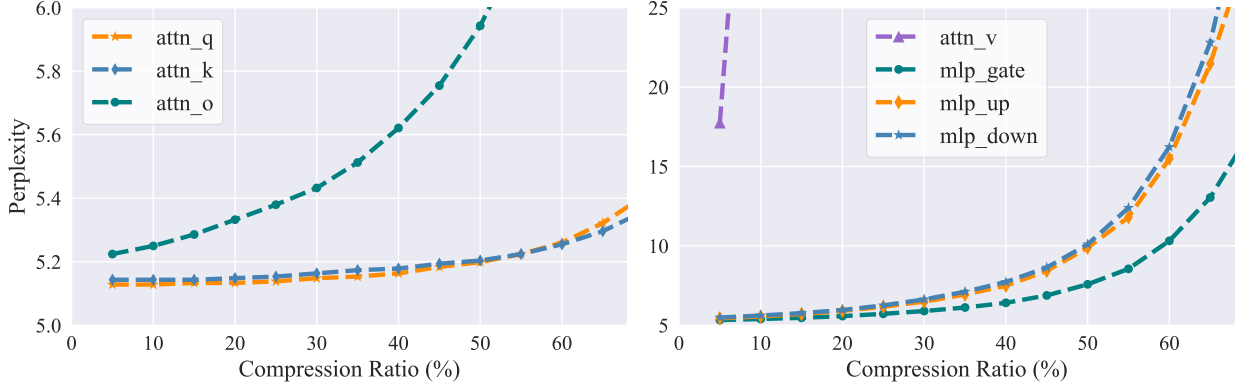


Figure 1: Sensitivity of different types of layers to low-rank compression. Each curve represents the compression of only that parameter type, with the horizontal axis indicating the compression ratio for that specific parameter type.

where $Cov(\mathbf{Y}) \in \mathbb{R}^{d_2 \times n}$, $E[\mathbf{Y}] \in \mathbb{R}^{d_2}$ is the covariance and mean of features. Thus, the original linear layer can be replaced by $\mathbf{B} = \mathbf{U}_r \in \mathbb{R}^{d_2 \times r}$, $\mathbf{A} = \mathbf{U}_r^T \mathbf{W} \in \mathbb{R}^{r \times d_1}$ and the bias compensation $\mathbf{b} = (\mathbf{I} - \mathbf{U}_r \mathbf{U}_r^T) E[\mathbf{Y}]$. We have observed that the current mainstream LLMs also exhibit characteristics of high-rank weights and low-rank features. Therefore, in this paper, we focus on the feature-based low-rank factorization.

2.2 Different Layers Exhibit Varying Degrees of Low-rank Sensitivity

Another challenge in LRC is allocating varying low-rank compression rates to different layers. Previous works have empirically or theoretically demonstrated that different components of Transformer-based masked language models and visual models exhibit distinct low-rank properties, such as the features of the self-attention modules having a lower rank compared to those of the feed-forward modules (Dong et al., 2023; Anagnostidis et al., 2022). These findings offer prior guidance for low-rank compression. However, detailed studies on current mainstream large language models are still lacking. Therefore, we take the LLaMA-v2-7b as an example to study the low-rank sensitivity within each layer across different types of layers and within the same type of layers. Llama-family LLMs have comprises seven distinct parameter categories: *attn_q*, *attn_k*, *attn_v*, *attn_o*, *mlp_up*, *mlp_down*, and *mlp_gate*. We evaluate the perplexity changes on Wikitext-2 (Merity et al., 2016) for each category under varying low-rank compression ratios. As Figure 1 shows, at the same low-rank compression rate, distinct layers exhibit notable performance variations. For *attn_q* and *attn_k*, they demonstrate robustness to low-rank compression,

with an increase in perplexity not exceeding 2% even at a compression rate of 60%. In contrast, *attn_v*, with an equivalent parameter count, exhibits high sensitivity, leading to a significant surge in perplexity with compression rates even below 5%. Therefore, assigning the same low-rank compression rate to different types of layers during low-rank compression of LLM is a sub-optimal solution. More empirical study results are shown in Appendix A.

3 Methodology

3.1 Feature-Based Low-Rank Decomposition in High-Dimensional Spaces

An efficient feature-based low-rank decomposition method performs PCA on features to identify the optimal low-rank matrices. To achieve general task-agnostic compression, we follow the setup of prior work (Frantar and Alistarh, 2023b; Sun et al., 2023; Ma et al., 2023), utilizing a subset of the pre-training data as calibration data $\mathcal{D}_{cal} = \{\mathbf{x}_i\}_{i=1}^n$. As described in Eq.4, we first estimate the covariance matrix of the entire feature space distribution \mathcal{Y} with the sample covariance matrix (SCM) of the calibration data features:

$$Cov_S(\mathbf{Y}) = \frac{1}{n-1} \sum_{i=1}^n (\mathbf{y}_i - \bar{\mathbf{y}})^T (\mathbf{y}_i - \bar{\mathbf{y}}), \quad (5)$$

where \mathbf{y}_i represents the feature of x_i , $\bar{\mathbf{y}}$ refers to the mean of all calibration data features. However, LLMs typically have high-dimensional feature spaces (e.g., the intermediate size of LLaMA-v2-7b has exceeded 10,000 dimensions). Precisely estimating the covariance matrix in such high-dimensional spaces has always been a statistical challenge, as the SCM does not effectively esti-

mate the covariance of high-dimensional distributions. For instance, calibration data sampled from pre-training datasets may introduce outlier features due to low-quality text or inadequate sampling. In high-dimensional spaces, these outlier features are difficult to identify due to the “curse of dimensionality”, and their impact is further exacerbated in estimating high-dimensional covariance matrices due to the sparsity of data points. Thus, to estimate the covariance of the feature space more robustly and accurately, we propose using the pooled covariance matrix (PCM) in place of the SCM. We split the calibration data into m groups. For each group, we can calculate the SCM $Cov_S(\mathbf{Y}_k)$, then the pooled covariance matrix is:

$$Cov_P(\mathbf{Y}) = \frac{1}{m} \sum_{k=1}^m Cov_S(\mathbf{Y}_k) \quad (6)$$

3.2 Low-Rank Allocation Based on Bayesian Optimization

As investigated in Section 2.2, different types of layers, and even each individual layer, exhibit varying sensitivities to low-rank compression. Therefore, allocating distinct compression ratios to different layers is crucial to achieve the desired compression rate with minimal performance degradation. For a LLM $f(\cdot; \theta)$, we compress it with the set of low-rank compression ratios $\lambda = \{\lambda_i\}_{i=1}^k$. We use a task-agnostic evaluation dataset \mathcal{D} to evaluate performance of the compressed model $f(\cdot; \theta, \lambda)$. For example, we can take the perplexity on a subset of pretraining data as the evaluation metric. Therefore, the optimization objective of low-rank allocation can be formulated as:

$$\begin{aligned} \min_{\lambda \in \mathcal{V}} H(\lambda) &= \mathbb{E}_{(x,y) \sim \mathcal{D}} h(f(x; \theta, \lambda), y) \\ s.t. \Sigma \lambda &\leq \rho, \end{aligned} \quad (7)$$

where $h(\cdot, \cdot)$ is the evaluation metric, ρ is model’s overall compression ratio. For LLMs, finding the optimal low-rank allocation based on a target compression ratio is a challenging optimization problem. First, the impact of the low-rank count allocated to different layers on the performance of the compressed model is combinatorial, and optimizing any one component independently may lead to a locally optimal solution. Second, due to LLMs’ vast number of parameters, evaluating $H(\lambda)$ is very time-consuming. Therefore, we leverage sample-efficient Bayesian optimization (Xu et al., 2022) to optimize Eq 7. Bayesian optimization (BO) estimates the objective $H(\lambda)$ using a

stochastic surrogate model and updates the posterior estimation of $H(\lambda)$ based on the results of each search step. We utilize the Gaussian process $\mathcal{N}(\mu(\cdot), \sigma^2(\cdot))$ as the surrogate model. Given the previous $t - 1$ search steps $\{\lambda_1, \dots, \lambda_{t-1}\}$ and their evaluation $H_{t-1} = [H(\lambda_1), \dots, H(\lambda_{t-1})]$, the surrogate model is updated as:

$$\begin{aligned} \mu(\lambda) &= \mathbf{k}(\mathbf{K} + \eta^2 \mathbf{I})^{-1} H_{t-1} \\ \sigma^2(\lambda) &= k(\lambda, \lambda) - \mathbf{k}^T (\mathbf{K} + \eta^2 \mathbf{I})^{-1} \mathbf{k}, \end{aligned} \quad (8)$$

where $k(\cdot, \cdot)$ is a kernel function, $\mathbf{k} = (k(\lambda, \lambda_i))_{i \in [t-1]}$, $\mathbf{K} = (k(\lambda_i, \lambda_j))_{i,j \in [t-1]}$, and $\eta^2 \mathbf{I}$ is the white kernel to model observation noise.

After obtaining the posterior estimation of $H(\lambda)$ (i.e., $H(\lambda) \sim \mathcal{N}(\mu(\lambda), \sigma^2(\lambda))$), BO determines the next compression rate allocation state through the acquisition function. Expected improvement (EI) is a popular and effective acquisition function:

$$\begin{aligned} \alpha(\lambda) &= \mathbb{E}_{H(\lambda)} [\max \{0, H' - H(\lambda)\}] \\ \lambda_t &= \underset{\lambda}{\operatorname{argmax}} \alpha(\lambda), \end{aligned} \quad (9)$$

where $H' = \min_{i \in [t-1]} H(\lambda_i)$, it means the minimal value observed so far. Then, BO chooses the point with the greatest expected improvement to explore. After obtaining the optimal ratio λ^* , we can determine the low rank. To fully leverage the acceleration effect of GPU matrix multiplication, we adhere to Nvidia’s user guidelines by rounding the low-rank dimensions to the nearest multiple of eight²:

$$r_i = \lceil \frac{(1 - \lambda_i) d_1 d_2}{d_1 + d_2} / 8 \rceil * 8. \quad (10)$$

During the Bayesian optimization process, to minimize the evaluation cost, we use a smaller validation set that may not comprehensively reflect the LLM’s performance, potentially leading to over-fitting in the optimization. To prevent BO from blindly improving the compressed model’s language modeling performance on the validation set, we aim for the compressed model to have a prediction distribution for the next word that is close to that of the original model. Hence, we employ the reverse KL divergence to quantify the difference:

$$\mathcal{L}(\theta, \lambda) = D_{KL}(f(x; \theta) || f(x; \theta, \lambda)). \quad (11)$$

²<https://docs.nvidia.com/deeplearning/performance/dl-performance-matrix-multiplication/index.html#gpu-imple>

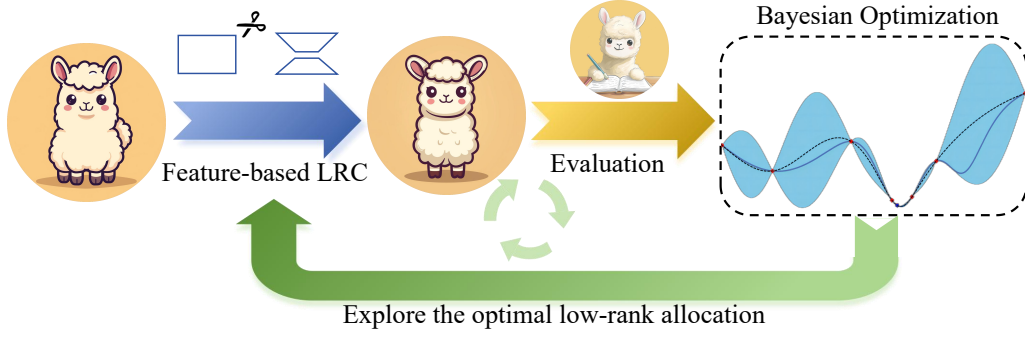


Figure 2: Illustration of our Bolaco. It initializes a low-rank dimension allocation and compresses the model via feature-based low-rank compression. Then, it evaluates the compression performance and optimizes the low-rank dimension allocation through Gaussian process-based Bayesian optimization.

3.3 Post-training

After low-rank compression, there remains a noticeable performance gap between the compressed model and the original LLM. To further bridge this gap, following Ma et al. (2023), we perform efficient low-rank subspace post-training on the compressed model. However, if we apply the original LoRA (Hu et al., 2022) to the low-rank compressed model, the tunable low-rank parameters may not be in the same subspace as the low-rank compressed model parameters, leading to an increase of the parameters’ rank after merging. Therefore, inspired by the ELoRA (Kopiczko et al., 2024), we select the subspace of compressed model parameters as fixed low-rank matrices and adjust the subspace by trainable vectors:

$$Y = (BA + \Lambda_b B_{r'} \Lambda_d A_{r'}) X, \quad (12)$$

where $B_{r'} \in \mathbb{R}^{d_2 \times r'}$ and $A_{r'} \in \mathbb{R}^{r' \times d_1}$ are fixed subspace of B and A , Λ_b and Λ_d are diagonal matrices. During the post-training, we only tune elements on the diagonal of Λ_b and Λ_d .

4 Experiments

4.1 Baseline Methods

We compare our method with the following competitive structured pruning and low-rank compression methods in LLMs.

LLM-Pruner (Ma et al., 2023) is a dependency-aware one-shot structured pruning. It evaluates the importance of each structure through a first-order Taylor expansion and prunes the structures with the lowest scores. After pruning, it uses LoRA post-training to recover performance.

FLAP (An et al., 2023) is an one-shot retraining-free structured pruning method. It utilizes a

fluctuation-based metric to measure the impact of pruning on features and employs a bias term to compensate for the pruning loss.

LoRD (Kaushal et al., 2023) is a naive feature-based low-rank compression method for code LLMs. It does not take into account the low-rank allocation of varying parameters. We migrate it to generic LLaMA-family LLMs.

4.2 Datasets

To evaluate the effectiveness of our proposed low-rank compression method in the task-agnostic setting, we conduct experiments in seven zero-shot common sense reasoning datasets: BoolQ (Clark et al., 2019), PIQA (Bisk et al., 2020), HellaSwag (Zellers et al., 2019), WinoGrande (Sakaguchi et al., 2021), ARC-easy/challenge (Clark et al., 2018) and OpenbookQA (Mihaylov et al., 2018). We also report the perplexity of the compressed model on the WikiText2 (Merity et al., 2016), PTB (Marcus et al., 1993), and C4 (Raffel et al., 2020) datasets to evaluate its language modeling capabilities.

4.3 Experimental Details

In our main experiments, we apply our method to LLaMA-v2-7b and LLaMA-v2-13b. We randomly select 1,024 samples from the training set of C4 as the calibration data. Each sample has a sequence length of 4,096. To estimate the covariance matrix while saving memory usage, we employ the Welford’s online algorithm (Welford, 1962). For the pooled covariance matrix, we partition the calibrated data into 32 groups. During the Bayesian optimization, we utilize the Matern kernel as the covariance function and use the test set of WikiText2 (about 34k tokens) as the evaluation data. Considering that Bayesian optimization is not well-suited for high-dimensional scenarios,




Figure 3: The perplexity of WikiText2 on LLaMA 2-7b with different compression ratios.

we conduct experiments with two settings based on the observations in Section 2.2: (a) 5×1 : We allow $attn_q$ and $attn_k$ to share a low-rank dimension, and the same type of parameters across different layers to also share a low-rank dimension, thus BO only optimizes 5 parameters; (b) 5×4 : Building on the setup of (a), we divide the model’s layers into 4 groups in sequence, with no parameter sharing between different groups, resulting in BO needing to optimize 20 parameters. We run 50 epochs of Bayesian optimization to search the optimal low-rank allocation. At the post-training stage, following LLM-Pruner, we use the Alpaca dataset (Taori et al., 2023) and train 2 epochs. More details can be found in the Appendix B.

4.4 Main Results

We report the perplexity of language modeling for various compression methods at different compression ratios in Figure 3 and 6, and the zero-shot common sense reasoning results in Table 1. In terms of language modeling capabilities, FLAP demonstrates strong competitiveness, particularly when the compression rate exceeds 30%, where FLAP’s perplexity is slightly better than our Bolaco (5×1). However, in the 7b model, Bolaco (5×4) achieves the best language modeling performance at high compression rates. Nevertheless, in the 13b model, despite Bolaco (5×4) still leading other compression techniques, it maintains a certain gap from FLAP. For zero-shot tasks, our method significantly outperforms all baselines without any further post-training, achieving an average performance increase of 1.5-2% across seven datasets. After post-training with only about 1% parameters and 3 hours, our method further narrows down the performance difference between the compressed

model and the original model. It retains 96%-98% of the original model’s performance at the 20% compression ratio, and at a 30% compression ratio, it maintains 91%-95% of the performance. Comparing the 5×1 and 5×4 setting, we find that the performance difference between the two is not significant. At the 20% compression ratio, simply allocating different low-rank dimensions to different types of parameters suffices to achieve the best current performance. However, at the 30% compression rate, the 5×4 setting outperforms the 5×1 , indicating that more granular low-rank assignments contribute to enhanced performance in compressed models at higher compression rates.

5 Analysis and Discussion

5.1 Impact of Calibration Data and Covariance Estimation

Accurate estimation of feature distribution is crucial for the feature-based low-rank decomposition, which primarily depends on the number of calibration samples and the accuracy of the covariance matrix estimation. Thus, we investigate the impact of the two factors on LLaMA-v2-7b at the 20% compression ratio. In this experiment, we do not account for the effects of low-rank dimensions allocation, and maintain consistency with the settings of LoRD. As results shown in Table 3, as the calibration dataset size gradually increases, we observe a consistent improvement in both the language modeling capabilities and the performance on downstream tasks of the compressed model. Therefore, given sufficient data and computational resources, expanding the calibration dataset is a reliable method for enhancing the performance of compressed models. On the other hand, comparing the two covariance estimation methods, there is no significant difference in their language modeling capabilities. However, for downstream common sense reasoning tasks, the pooled SCM achieves an average improvement of 0.3 points across seven datasets without any additional burden.

5.2 Impact of Objective Function

We explore the impact of the objective function in the Bayesian optimization stage. We conduct experiments on LLaMA-v2-7b and report results in Table 3. Overall, incorporating the reverse KL divergence (RKL) between the compressed model and the original model’s predictive distribution into the objective function can lead to a better low-rank di-

Ratio	Methods	BoolQ	PIQA	HellaSwag	WinoGrande	ARC-e	ARC-c	OBQA	Average
0%	LLaMA-v2-7b	77.74	78.07	75.97	68.98	76.30	46.33	44.20	66.80
20%	LLM-Pruner	63.27	76.12	67.93	64.80	68.73	38.65	40.00	59.93
	LLM-Pruner (w/ PT)	66.45	76.28	70.90	65.75	70.62	39.59	43.20	61.83
	FLAP	70.21	75.24	69.34	66.30	67.30	39.42	37.40	60.74
	LoRD	72.60	73.56	63.70	65.90	69.70	37.71	39.20	60.34
	Bolaco (5×1)	71.44	75.24	67.22	67.09	72.85	39.16	42.00	62.14
	Bolaco (5×1 w/ PT)	<u>74.56</u>	76.61	<u>71.74</u>	<u>66.77</u>	73.57	42.75	43.20	64.17
	Bolaco (5×4)	71.44	74.70	67.24	66.54	71.84	38.31	<u>42.60</u>	61.81
	Bolaco (5×4 w/ PT)	75.90	<u>76.44</u>	71.85	67.09	<u>72.98</u>	<u>41.89</u>	42.20	<u>64.05</u>
30%	LLM-Pruner	52.51	71.93	59.49	58.72	61.41	33.96	36.60	53.52
	LLM-Pruner (w/ PT)	63.30	76.01	65.23	64.25	66.62	37.20	40.20	58.97
	FLAP	66.88	72.74	63.80	64.01	60.65	34.47	36.40	56.99
	LoRD	69.63	70.46	55.87	64.17	63.80	32.59	35.00	55.93
	Bolaco (5×1)	69.11	71.55	56.02	63.69	64.52	33.02	38.40	56.62
	Bolaco (5×1 w/ PT)	70.80	<u>74.81</u>	<u>67.38</u>	<u>64.64</u>	<u>68.35</u>	38.65	41.00	<u>60.80</u>
	Bolaco (5×4)	73.85	71.27	57.85	64.48	65.87	34.30	37.80	57.92
	Bolaco (5×4 w/ PT)	<u>71.01</u>	74.59	67.68	65.75	69.07	38.99	<u>40.80</u>	61.13
0%	LLaMA-v2-13b	80.52	79.05	79.38	72.14	79.42	49.23	45.20	69.27
20%	LLM-Pruner	66.33	78.18	74.47	64.48	72.26	<u>45.90</u>	44.20	63.69
	LLM-Pruner (w/ PT)	67.06	78.94	75.92	67.32	72.69	44.28	44.60	64.40
	FLAP	71.28	76.55	74.67	69.53	72.56	44.03	42.00	64.37
	LoRD	78.47	76.01	69.58	<u>71.03</u>	74.33	40.87	44.40	64.96
	Bolaco (5×1)	80.58	77.04	73.11	70.01	<u>76.39</u>	44.28	45.00	66.63
	Bolaco (5×1 w/ PT)	81.25	<u>78.51</u>	76.70	71.35	77.10	47.44	44.80	68.16
	Bolaco (5×4)	78.81	76.33	72.50	69.53	75.42	42.23	45.00	65.69
	Bolaco (5×4 w/ PT)	79.33	78.07	<u>76.41</u>	70.80	75.29	44.80	<u>44.80</u>	<u>67.07</u>
30%	LLM-Pruner	62.45	75.90	67.90	60.22	65.45	40.36	<u>44.60</u>	59.55
	LLM-Pruner (w/ PT)	68.29	76.66	72.03	64.09	69.20	41.13	45.40	62.40
	FLAP	65.54	74.81	70.29	67.48	67.38	38.23	40.00	60.53
	LoRD	75.05	73.88	63.08	69.46	69.78	39.16	38.60	61.29
	Bolaco (5×1)	76.42	74.86	64.36	67.01	71.93	40.19	40.60	62.20
	Bolaco (5×1 w/ PT)	<u>79.39</u>	<u>76.33</u>	<u>73.21</u>	69.46	74.12	<u>43.77</u>	44.20	<u>65.78</u>
	Bolaco (5×4)	75.63	74.16	66.66	<u>69.06</u>	71.76	40.10	41.60	62.71
	Bolaco (5×4 w/ PT)	80.34	75.79	73.98	67.88	<u>73.40</u>	45.14	<u>44.60</u>	65.88

Table 1: Zero-shot performance of the compressed LLaMA-v2 models. w/ PT means the method with post-training. **Bold** denotes the best result at the same compression ratio, while underline indicates the second best result.

	Wikitext (\downarrow)	PTB (\downarrow)	C4 (\downarrow)	ZS (\uparrow)
<i>Covariance estimate</i>				
Naive SCM	9.96	54.69	11.46	60.34
Pooled SCM	9.93	54.68	11.45	60.64
<i># Samples</i>				
128	10.55	56.29	11.99	60.26
256	10.24	55.42	11.88	60.16
512	10.30	55.03	11.61	60.56
1,024	9.93	54.68	11.45	60.64

Table 2: Impact of different covariance estimation methods and the number of calibration data. “ZS” denotes the average performance on seven zero-shot common sense reasoning datasets.

term may serve two roles. Firstly, as a regularization term, it prevents overfitting on smaller validation sets during Bayesian optimization. Although the compressed model exhibits a slight increase in perplexity on the language modeling dataset at the 20% compression rate with the 5×4 setting, there is a significant improvement in performance on downstream tasks. Secondly, incorporating the RKL term may smooth the objective function, enabling the Gaussian process surrogate model to more accurately approximate the real black-box objective function.

5.3 The Transferability of Low-rank Allocation

In practical applications, we may utilize a variety of fine-tuned models based on the LLaMA founda-

mensions allocation. Especially in the 5×4 setting, which is more difficult to optimize for Bayesian optimization, the performance gains from RKL term are even more obvious. We suppose that the RKL

	Wikitext (\downarrow)	PTB (\downarrow)	Zero-shot (\uparrow)
20%			
PPL (5×1)	8.04	48.94	61.83
w/ RKL (5×1)	7.61	48.37	62.14
PPL (5×4)	7.78	48.37	61.05
w/ RKL (5×4)	8.08	48.85	61.81
30%			
PPL (5×1)	13.80	76.70	56.84
w/ RKL (5×1)	13.91	81.64	56.62
PPL (5×4)	12.74	77.30	56.32
w/ RKL (5×4)	11.99	71.49	57.92

Table 3: Results under different objective function. “Zero-shot” denotes the average performance on seven zero-shot common sense reasoning datasets.




Figure 4: The average performance on zero-shot tasks.

tion model. If we perform Bayesian optimization from scratch to optimize the low-rank allocation for each model, it will waste a significant amount of time and computational resources. Hence, we investigate whether the low-rank allocation of the base model can be transferred to the corresponding fine-tuned models. We transfer the allocation of LLaMA-v2-7b/13b to LLaMA-v2-7b/13b-chat, respectively. We consider two migration strategies: **a)** directly reusing the low-rank allocation of the base model and **b)** using the low-rank allocation of the base model as the initial value for Bayesian optimization and then optimizing only 20 epochs. As Figure 4 shows, direct reuse can achieve results that outperform all baseline methods, even the Bayesian optimization from scratch. If 20 epochs of Bayesian optimization follow reuse, there is a chance to find an even better low-rank allocation.

6 Related work

A common technique for low-rank factorization is SVD, which retains only the top- r largest singular values and their corresponding singular vectors to obtain two rank- r matrices. Ben Noach and Goldberg (2020) first combine SVD with knowl-

edge distillation, applying it to compress BERT. Directly applying SVD decomposition implies an assumption that each parameter in the weight matrix equally affects the model performance. This contradicts many previous research, therefore, FWSVD (Hsu et al., 2022) and TFWSVD (Hua et al., 2022) consider weighting the weight matrix using Fisher information. Chen et al. (2021) observe that PLMs’ weights are not inherently low-rank matrices. Therefore, directly applying SVD will result in significant reconstruction loss. However, they find that the product of data representation and weights is low-rank. Hence, they perform a global low-rank decomposition on it. Following this observation, Yu and Wu (2023) propose the atomic feature mimicking (AFM) method to decompose the output features. Ren and Zhu (2023) also observe the high rank phenomenon of PLM weights. They utilize iterative first-order unstructured pruning to reduce the rank of the weight matrix, and then apply Fisher information-weighted SVD decomposition for low-rank compression. For LLMs, low-rank compression has not yet received the attention it deserves. LoRD (Kaushal et al., 2023) applies AFM to code LLMs, demonstrating the potential of low-rank decomposition in compressing LLM. Recently, Sharma et al. (2023) conduct an in-depth study on the weight decomposition of LLMs and discover that the low-rank components of the weights encapsulate low-frequency information. By meticulously selecting low-rank components, it is possible to eliminate interfering signals and further improve LLMs’ performance. However, their research does not propose a practical low-rank compression algorithm.

7 Conclusion

In this paper, we attempt to unearth the potential of low-rank compression for lightweight universal LLMs. We thoroughly investigate the challenges of low-rank compression in LLMs and the low-rank characteristics of features within LLMs. We propose a Bayesian optimization-based feature low-rank compression to address these challenges, incorporating pooled covariance estimation and Bayesian optimization for more precise feature distribution estimation and low-rank dimension allocation, respectively. Experimental results on the LLaMA 2 model demonstrate that our method significantly outperforms existing structured pruning and other low-rank compression techniques.

Limitations

Although our proposed Bolaco has made significant progress in low-rank compression for LLMs, there are still some limitations:

- Due to computational resource constraints, we only conduct thorough experiments on two commonly used LLaMA 2 models, lacking investigation into larger models (such as LLaMA 2-70B), other architectures (such as the OPT and T5 families), and multimodal models.
- To improve the efficiency of Bayesian optimization, we reduced the parameter dimensions by sharing parameters of different types and layers using low-rank dimensions. This may limit the potential performance of the model. We plan to use more advanced methods to find better low-rank allocation while maintaining flexibility.
- Compared to state-of-the-art structured pruning, low-rank compression falls short in highly compressed language modeling but exhibits better zero-shot performance on downstream tasks. These observations inspire us to investigate how to effectively combine these two approaches to capitalize on their advantages in the future.

References

- Yongqi An, Xu Zhao, Tao Yu, Ming Tang, and Jinqiao Wang. 2023. [Fluctuation-based adaptive structured pruning for large language models](#).
- Sotiris Anagnostidis, Luca Biggio, Lorenzo Noci, Antonio Orvieto, Sidak Pal Singh, and Aurelien Lucchi. 2022. [Signal propagation in transformers: Theoretical perspectives and the role of rank collapse](#). In *Advances in Neural Information Processing Systems*.
- Matan Ben Noach and Yoav Goldberg. 2020. [Compressing pre-trained language models by matrix decomposition](#). In *Proceedings of the 1st Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 10th International Joint Conference on Natural Language Processing*, pages 884–889, Suzhou, China. Association for Computational Linguistics.
- Yonatan Bisk, Rowan Zellers, Ronan bras, Jianfeng Gao, and Choi Yejin. 2020. [Piqa: Reasoning about physical commonsense in natural language](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 34:7432–7439.
- Sébastien Bubeck, Varun Chandrasekaran, Ronen Eldan, Johannes Gehrke, Eric Horvitz, Ece Kamar, Peter Lee, Yin Tat Lee, Yuanzhi Li, Scott Lundberg, Harsha Nori, Hamid Palangi, Marco Tulio Ribeiro, and Yi Zhang. 2023. [Sparks of artificial general intelligence: Early experiments with gpt-4](#).
- Patrick Chen, Hsiang-Fu Yu, Inderjit Dhillon, and Chojui Hsieh. 2021. [Drone: Data-aware low-rank compression for large nlp models](#). In *Advances in Neural Information Processing Systems*, volume 34, pages 29321–29334. Curran Associates, Inc.
- Christopher Clark, Kenton Lee, Ming-Wei Chang, Tom Kwiatkowski, Michael Collins, and Kristina Toutanova. 2019. [BoolQ: Exploring the surprising difficulty of natural yes/no questions](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 2924–2936, Minneapolis, Minnesota. Association for Computational Linguistics.
- 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](#).
- Yihe Dong, Jean-Baptiste Cordonnier, and Andreas Loukas. 2023. [Attention is not all you need: Pure attention loses rank doubly exponentially with depth](#).
- Elias Frantar and Dan Alistarh. 2023a. [SparseGPT: Massive language models can be accurately pruned in one-shot](#). In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 10323–10337. PMLR.
- Elias Frantar and Dan Alistarh. 2023b. [SparseGPT: Massive language models can be accurately pruned in one-shot](#). *arXiv preprint arXiv:2301.00774*.
- Elias Frantar, Saleh Ashkboos, Torsten Hoefer, and Dan Alistarh. 2023. [OPTQ: Accurate quantization for generative pre-trained transformers](#). In *The Eleventh International Conference on Learning Representations*.
- Yuxian Gu, Li Dong, Furu Wei, and Minlie Huang. 2024. [MiniLLM: Knowledge distillation of large language models](#). In *The Twelfth International Conference on Learning Representations*.
- Jordan Hoffmann, Sebastian Borgeaud, Arthur Mensch, Elena Buchatskaya, Trevor Cai, Eliza Rutherford, Diego de Las Casas, Lisa Anne Hendricks, Johannes Welbl, Aidan Clark, Tom Hennigan, Eric Noland, Katie Millican, George van den Driessche, Bogdan Damoc, Aurelia Guy, Simon Osindero, Karen Simonyan, Erich Elsen, Jack W. Rae, Oriol Vinyals, and Laurent Sifre. 2022. [Training compute-optimal large language models](#).
- Yen-Chang Hsu, Ting Hua, Sungen Chang, Qian Lou, Yilin Shen, and Hongxia Jin. 2022. [Language model compression with weighted low-rank factorization](#). In *International Conference on Learning Representations*.

681	Edward J Hu, yelong shen, Phillip Wallis, Zeyuan Allen-	<i>Empirical Methods in Natural Language Processing</i> ,	736
682	Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu	pages 2381–2391, Brussels, Belgium. Association	737
683	Chen. 2022. LoRA: Low-rank adaptation of large	for Computational Linguistics.	738
684	language models . In <i>International Conference on</i>		
685	<i>Learning Representations</i> .	OpenAI. 2022. Chatgpt: Optimizing language models	739
		for dialogue. <i>Open AI, blog</i> .	740
686	Ting Hua, Yen-Chang Hsu, Felicity Wang, Qian Lou,	OpenAI. 2023. Gpt-4 technical report. <i>ArXiv</i> ,	741
687	Yilin Shen, and Hongxia Jin. 2022. Numerical opti-	abs/2303.08774.	742
688	mizations for weighted low-rank estimation on lan-		
689	guage models . In <i>Proceedings of the 2022 Confer-</i>	Colin Raffel, Noam Shazeer, Adam Roberts, Kather-	743
690	<i>ence on Empirical Methods in Natural Language Pro-</i>	ine Lee, Sharan Narang, Michael Matena, Yanqi	744
691	<i>cessing</i> , pages 1404–1416, Abu Dhabi, United Arab	Zhou, Wei Li, and Peter J. Liu. 2020. Exploring the	745
692	Emirates. Association for Computational Linguistics.	limits of transfer learning with a unified text-to-text	746
		transformer . <i>Journal of Machine Learning Research</i> ,	747
693	Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B.	21(140):1–67.	748
694	Brown, Benjamin Chess, Rewon Child, Scott Gray,		
695	Alec Radford, Jeffrey Wu, and Dario Amodei. 2020.	Elias Raninen, David E. Tyler, and Esa Ollila. 2022.	749
696	Scaling laws for neural language models .	Linear pooling of sample covariance matrices . <i>IEEE</i>	750
		<i>Transactions on Signal Processing</i> , 70:659–672.	751
697	Ayush Kaushal, Tejas Vaidhya, and Irina Rish. 2023.	Siyu Ren and Kenny Q. Zhu. 2023. Low-rank prune-	752
698	Lord: Low rank decomposition of monolingual code	and-factorize for language model compression .	753
699	llms for one-shot compression .		
700	Dawid J. Kopiczko, Tijmen Blankevoort, and Yuki M.	Keisuke Sakaguchi, Ronan Le Bras, Chandra Bhagavat-	754
701	Asano. 2024. ELoRA: Efficient low-rank adaptation	ula, and Yejin Choi. 2021. Winogrande: an adver-	755
702	with random matrices . In <i>The Twelfth International</i>	sarial winograd schema challenge at scale . <i>Commun.</i>	756
703	<i>Conference on Learning Representations</i> .	<i>ACM</i> , 64(9):99–106.	757
704	Yixiao Li, Yifan Yu, Qingru Zhang, Chen Liang,	Pratyusha Sharma, Jordan T. Ash, and Dipendra Misra.	758
705	Pengcheng He, Weizhu Chen, and Tuo Zhao. 2023.	2023. The truth is in there: Improving reasoning in	759
706	LoSparse: Structured compression of large language	language models with layer-selective rank reduction .	760
707	models based on low-rank and sparse approximation .		
708	In <i>Proceedings of the 40th International Conference</i>	Mingjie Sun, Zhuang Liu, Anna Bair, and J. Zico	761
709	<i>on Machine Learning</i> , volume 202 of <i>Proceedings</i>	Kolter. 2023. A simple and effective pruning ap-	762
710	<i>of Machine Learning Research</i> , pages 20336–20350.	proach for large language models. <i>arXiv preprint</i>	763
711	PMLR.	<i>arXiv:2306.11695</i> .	764
712	Ji Lin, Jiaming Tang, Haotian Tang, Shang Yang,	Mingjie Sun, Zhuang Liu, Anna Bair, and J Zico Kolter.	765
713	Xingyu Dang, Chuang Gan, and Song Han. 2023.	2024. A simple and effective pruning approach for	766
714	Awq: Activation-aware weight quantization for llm	large language models . In <i>The Twelfth International</i>	767
715	compression and acceleration .	<i>Conference on Learning Representations</i> .	768
716	Zechun Liu, Barlas Oguz, Changsheng Zhao, Ernie	Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann	769
717	Chang, Pierre Stock, Yashar Mehdad, Yangyang Shi,	Dubois, Xuechen Li, Carlos Guestrin, Percy Liang,	770
718	Raghuraman Krishnamoorthi, and Vikas Chandra.	and Tatsunori B. Hashimoto. 2023. Stanford alpaca:	771
719	2023. LLM-qat: Data-free quantization aware training	An instruction-following llama model. https://	772
720	for large language models .	github.com/tatsu-lab/stanford_alpaca .	773
721	Xinyin Ma, Gongfan Fang, and Xinchao Wang. 2023.	Hugo Touvron, Louis Martin, Kevin Stone, Peter Al-	774
722	LLM-pruner: On the structural pruning of large lan-	bert, Amjad Almahairi, Yasmine Babaei, Nikolay	775
723	guage models . In <i>Thirty-seventh Conference on Neu-</i>	Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti	776
724	<i>ral Information Processing Systems</i> .	Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton	777
725	Mitchell P. Marcus, Beatrice Santorini, and Mary Ann	Ferrer, Moya Chen, Guillem Cucurull, David Esiobu,	778
726	Marcinkiewicz. 1993. Building a large annotated cor-	Jude Fernandes, Jeremy Fu, Wenxin Fu, Brian Fuller,	779
727	pus of English: The Penn Treebank . <i>Computational</i>	Cynthia Gao, Vedanuj Goswami, Naman Goyal, An-	780
728	<i>Linguistics</i> , 19(2):313–330.	thony Hartshorn, Saghar Hosseini, Rui Hou, Hakan	781
729	Stephen Merity, Caiming Xiong, James Bradbury, and	Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa,	782
730	Richard Socher. 2016. Pointer sentinel mixture mod-	Isabel Kloumann, Artem Korenev, Punit Singh Koura,	783
731	els .	Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Di-	784
732	Todor Mihaylov, Peter Clark, Tushar Khot, and Ashish	ana Liskovich, Yinghai Lu, Yuning Mao, Xavier Mar-	785
733	Sabharwal. 2018. Can a suit of armor conduct elec-	tinnet, Todor Mihaylov, Pushkar Mishra, Igor Moly-	786
734	tricity? a new dataset for open book question an-	bog, Yixin Nie, Andrew Poulton, Jeremy Reizen-	787
735	swering . In <i>Proceedings of the 2018 Conference on</i>	stein, Rashi Rungra, Kalyan Saladi, Alan Schelten,	788
		Ruan Silva, Eric Michael Smith, Ranjan Subrama-	789
		nian, Xiaoqing Ellen Tan, Binh Tang, Ross Tay-	790
		lor, Adina Williams, Jian Xiang Kuan, Puxin Xu,	791

Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurelien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. 2023. [Llama 2: Open foundation and fine-tuned chat models](#).

Zhongwei Wan, Xin Wang, Che Liu, Samiul Alam, Yu Zheng, Jiachen Liu, Zhongnan Qu, Shen Yan, Yi Zhu, Quanlu Zhang, Mosharaf Chowdhury, and Mi Zhang. 2024. [Efficient large language models: A survey](#).

B. P. Welford. 1962. [Note on a method for calculating corrected sums of squares and products](#). *Technometrics*, 4(3):419–420.

Wenjie Xu, Yuning Jiang, Emilio T. Maddalena, and Colin N. Jones. 2022. [Lower bounds on the worst-case complexity of efficient global optimization](#).

Zhengyuan Yang, Linjie Li, Kevin Lin, Jianfeng Wang, Chung-Ching Lin, Zicheng Liu, and Lijuan Wang. 2023. [The dawn of lmms: Preliminary explorations with gpt-4v\(ision\)](#).

Hao Yu and Jianxin Wu. 2023. [Compressing transformers: Features are low-rank, but weights are not!](#) *Proceedings of the AAAI Conference on Artificial Intelligence*, 37(9):11007–11015.

Rowan Zellers, Ari Holtzman, Yonatan Bisk, Ali Farhadi, and Yejin Choi. 2019. [HellaSwag: Can a machine really finish your sentence?](#) In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4791–4800, Florence, Italy. Association for Computational Linguistics.

A More Experiments on Low-rank Sensitivity

As shown in Figure 5. We further reduce the parameters by 50% on LLaMA-v2-7b by low-rank compression for each layer and test its perplexity on WikiText2. We observe that the low-rank sensitivity varies significantly across different types of parameters. Compression of $attn_q$ and $attn_k$ seemingly has negligible impact on overall performance across all layers. In contrast, the upper layers of mlp are more sensitive compared to the lower and middle layers. We also conduct experiments on LLaMA-7b-chat and OPT-6.7b and find significant variability between all the different types of parameters. However, for OPT-6.7b, the differences between them are less pronounced than for LLama, especially for $attn_v$, which does not show an explosive increase in perplexity.

B Implementation Details

For LoRD, due to the absence of reference settings for its application on the LLaMA, we manually search a good low-rank allocation for it. At the 20% compression ratio, we do not compress $attn_v$, and reduce the parameter count of $attn_q/k$ by 30%, with a 20% reduction in the remaining parameters. At the 30% ratio, we reduce the parameter count of $attn_q/k$ by 45%, with a 30% reduction in the remaining parameters except $attn_v$.

At the post-training stage, we only add fine-tunable low-rank matrices for the compressed parameters. We set the low-rank dimension $r' = 256$, the learning rate is $1e-2$, and the batch size is 64.

C Statistics of the Compressed Model

We report the statistic of original and compressed models in Table 4, including the parameter count, MACs and memory requirements. The statistical evaluation is conducted using the inference mode, where the model is fed a sentence consisting of 64 tokens.

To aid subsequent researchers in reproducing our results, Table 5 provides the low-rank allocations of Bolaco. The elements of the array represent the low-rank dimensions for $attn_q/k$, $attn_o$, mlp_gate , mlp_up , and mlp_down , respectively. ‘NA’ denotes that the parameter is not compressed.

Method	Ratio	#Params	MACs	Memory
LLaMA 2-7b	0%	6.74B	423.98G	12.62GiB
LLM-Pruner	20%	5.42B	340.48G	10.16GiB
FLAP	20%	5.45B	342.30G	10.22GiB
LoRD	20%	5.45B	370.12G	10.32GiB
Bolaco (5×1)	20%	5.44B	388.95G	10.28GiB
Bolaco (5×4)	20%	5.44B	391.18G	10.25GiB
LLM-Pruner	30%	4.84B	302.83G	9.17GiB
FLAP	30%	4.80B	300.72G	9.04GiB
LoRD	30%	4.79B	341.91G	9.07GiB
Bolaco (5×1)	30%	4.79B	359.48G	9.04GiB
Bolaco (5×4)	30%	4.80B	356.03G	9.06GiB
LLaMA 2-13b	0%	13.02B	824.26G	24.45GiB
LLM-Pruner	20%	10.48B	662.95G	19.75GiB
FLAP	20%	10.48B	663.85G	19.64GiB
LoRD	20%	10.49B	717.86G	19.79GiB
Bolaco (5×1)	20%	10.48B	777.58G	19.71GiB
Bolaco (5×4)	20%	10.48B	772.16G	19.69GiB
LLM-Pruner	30%	9.21B	581.40G	17.35GiB
FLAP	30%	9.21B	582.72G	17.29GiB
LoRD	30%	9.21B	663.15G	17.38GiB
Bolaco (5×1)	30%	9.21B	708.16G	17.36GiB
Bolaco (5×4)	30%	9.21B	694.58G	17.35GiB

Table 4: Statistics of the compressed model.

Model	Method	Ratio	Low rank allocation
LLaMA-v2-7b	Bolaco (5×1)	20%	[824, 1696, 2208, NA, 2592]
	Bolaco (5×4)	20%	[[992, 1640, NA, NA, NA], [824, 1808, 2360, NA, 2352], [976, NA, 2696, 2400, 2736], [408, 1000, 2944, 2656, 1304]]
	Bolaco (5×1)	30%	[624, 1368, 1760, 2592, 2464]
	Bolaco (5×4)	30%	[[736, 1344, 2568, 2704, NA], [904, 1904, 1304, 2600, 2376], [632, 1456, 2080, 1528, 2304], [528, 1040, 2248, 2376, 1776]]
LLaMA-v2-13b	Bolaco (5×1)	20%	[632, 1400, 2184, NA, 2976]
	Bolaco (5×4)	20%	[[816, 1536, 2512, 2984, NA], [408, 1496, 2696, NA, 2960], [1048, 928, 2080, 2552, 2288], [968, 840, NA, 2560, 2648]]
	Bolaco (5×1)	30%	[520, 1256, 1824, 2616, 2560]
	Bolaco (5×4)	30%	[[600, 1256, 2568, 2744, NA], [920, 1472, 2464, 2520, 2120], [848, 1208, 2048, 1584, 2360], [696, 816, 1928, 2328, 1448]]

Table 5: The low-rank allocation of our Bolaco.

D Language Modeling Capabilities for Compressed Models

Figure 6 illustrates the perplexity changes on WikiText, PTB, and C4 datasets for different compression methods on LLaMA-v2-7b and 13b as the compression rate increases.

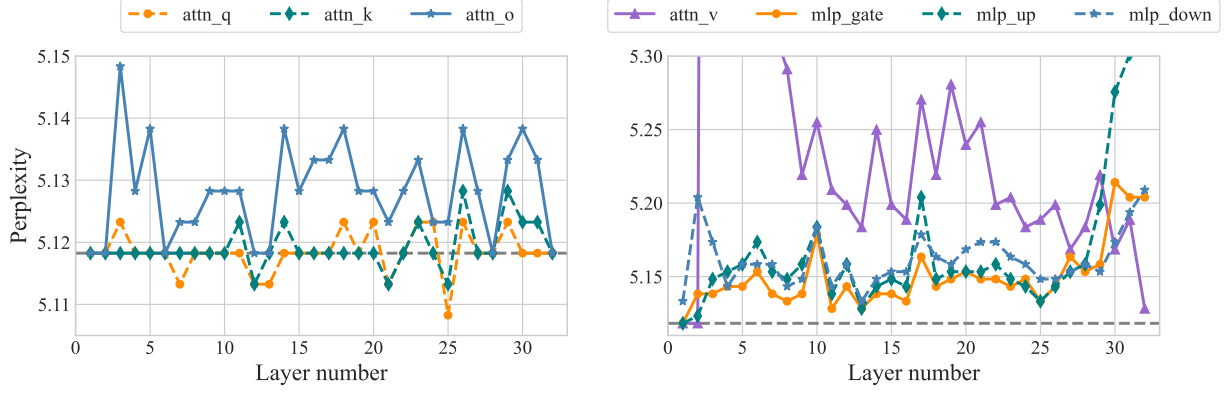
E Case Study

We showcase the generation results of the LLaMA-v2-7b and its compression model via Bolaco in Table 6. We observe that models compressed via Bo-

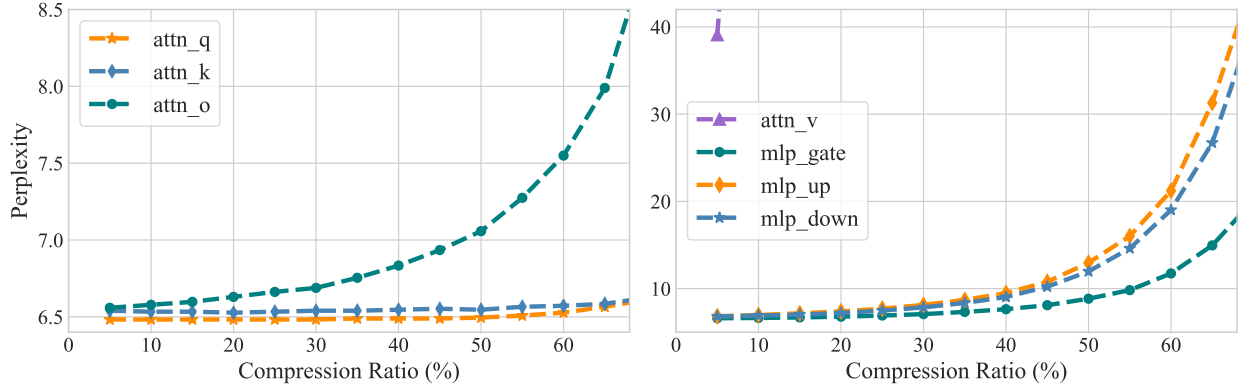
Model	Generated Sentences
LLaMA 2-7b	Instruction: I sit in front of a computer all day. How do I manage and mitigate eye strain? Response: Eye strain is a common complaint among computer users. It is caused by the eye's inability to focus on the computer screen. This is due to the fact that the computer screen is too close to the eyes. The eyes are also strained by the glare from the computer screen. The best way to manage and mitigate eye strain is to take frequent breaks from the computer. This will allow the eyes to rest and relax.
	Instruction: Translate the sentence 'I have no mouth but I must scream' into Spanish. Response: No tengo boca pero debo gritar.
	Instruction: Who is Donald Trump? Response: Donald Trump is the 45th President of the United States.
Bolaco (20%) (w/o tune)	Instruction: I sit in front of a computer all day. How do I manage and mitigate eye strain? Response: I have a few tips for you. First, I recommend that you take frequent breaks to rest your eyes. I also recommend that you take frequent breaks to rest your eyes.
	Instruction: Translate the sentence 'I have no mouth but I must scream' into Spanish. Response: I have no mouth but I must scream.
	Instruction: Who is Donald Trump? Response: Donald Trump is the president of the United States.
Bolaco (20%) (w/ tune)	Instruction: I sit in front of a computer all day. How do I manage and mitigate eye strain? Response: Eye strain is a common problem that can occur when the eyes are exposed to prolonged periods of bright light, such as from a computer screen. To mitigate eye strain, it is important to take regular breaks and look away from the screen every 20 minutes. Additionally, it is recommended to use a computer with a low glare monitor, adjust the brightness of the screen, and use a headset or earphones to reduce the amount of noise from the
	Instruction: Translate the sentence 'I have no mouth but I must scream' into Spanish. Response: Yo no tengo boca pero debo gritar.
	Instruction: Who is Donald Trump? Response: Donald Trump is the 45th President of the United States of America. He was born on December 4, 1942, in New York City. He is a businessman, television personality, and politician. He is the founder and owner of the Trump Organization, which is a real estate development and management company. He is also the former CEO of the Trump Entertainment Resorts, which operates casinos and hotels.

Table 6: Generated Examples from LLaMA-v2-7b and Bolaco.


laco tend to produce brief and repetitive responses to prompts without post-training. However, this issue can be resolved after efficient post-training, resulting in smooth and informative replies.



(a) Low rank sensitivity of individual layer on LLaMA-v2-7b.




(b) Sensitivity of different types of layers to low-rank compression on the LLaMA-v2-7b-chat.




(c) Sensitivity of different types of layers to low-rank compression on the OPT-6.7b.


Figure 5: More results on low-rank sensitivity.



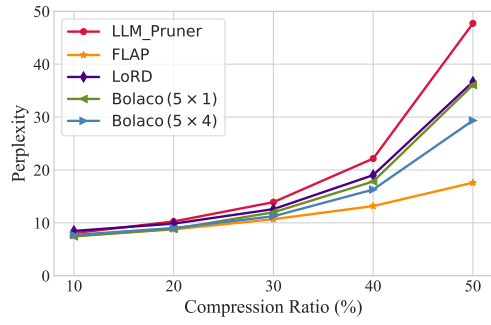
(a) The perplexity of WikiText2 on LLaMA-v2-7b




(b) The perplexity of WikiText2 on LLaMA-v2-13b




(c) The perplexity of C4 on LLaMA-v2-7b



(d) The perplexity of C4 on LLaMA-v2-13b



(e) The perplexity of PTB on LLaMA-v2-7b



(f) The perplexity of PTB on LLaMA-v2-13b

Figure 6: Language modeling capabilities at different compression ratios.