

ASER: Activation Smoothing and Error Reconstruction for Large Language Model Quantization

Weibo Zhao*, Yubin Shi*[†], Xinyu Lyu, Wanchen Sui, Shen Li, Yong Li

Alibaba Cloud Computing
{weibo.zwb, shiyubin.syb}@alibaba-inc.com

Abstract

Quantization stands as a pivotal technique for deploying large language models (LLMs), yet it poses significant challenges particularly in achieving effective low-bit quantization. The limited numerical mapping makes the quantized model produce a non-trivial error, bringing out intolerable performance degradation. This paper is anchored in the basic idea of model compression objectives, and delves into the layer-wise error distribution of LLMs during post-training quantization. Subsequently, we introduce ASER, an algorithm consisting of (1) **Error Reconstruction**: low-rank compensation for quantization error with LoRA-style matrices constructed by whitening SVD; (2) **Activation Smoothing**: outlier extraction to gain smooth activation and better error compensation. ASER is capable of quantizing typical LLMs to low-bit ones, particularly preserving accuracy even in W4A8 per-channel setup. Experimental results show that ASER is competitive among the state-of-the-art quantization algorithms, showing potential to activation quantization, with minor overhead.

Introduction

The proliferation of large language models (LLMs) places increasing demands on computation and storage (Vaswani et al. 2017; Devlin et al. 2018), and stretches the boundaries of contemporary hardware prowess. Regarding typical generative models such as Llama3.1-310B (Llama Team 2024) or Qwen2-72B (Yang et al. 2024), they require hundreds of gigabytes of VRAM, and often rely on multi-GPU cluster data centers for deployment. Researchers have been trying to compress these large models while maintaining their performance in a more compact form. Quantization (Frantar et al. 2022; Xiao et al. 2023) of large language models emerges as a pivotal strategy in optimizing the deployment and computational efficiency of these heavy neural architectures.

Quantization involves converting the model’s weights and activations from floating-point numbers to lower-precision integers, whose primary goal is alleviating the computational and memory requirements without significantly compromising the model’s performance, enabling edge deployment. Take post training quantization (PTQ) as an example,

current methods try to smooth (Xiao et al. 2023), rotate (Liu et al. 2024) or transform representations in models, showing promise of achieving quantization-friendly data range and remarkable compression rate. However, when it comes to low-bit quantization, even with the application of the aforementioned techniques, the error introduced by the quantized model becomes evident, resulting in notable degradation.

For typical quantization, the optimization objective can be formulated as minimizing $\|\mathbf{W}\mathbf{X} - \mathbf{W}_q\mathbf{X}\|_F$, where \mathbf{W} and $\mathbf{W}_q = Q(\mathbf{W})$ are original model weight and its quantized one respectively. We empirically find this inevitable quantization error has low-rank property, whose singular value distribution featuring a small number of high values and a long-tail bulk of low ones. This varies both intra and inter Transformer blocks, which prompts us to consider using a module similar to LoRA (Hu et al. 2021) to reconstruct this error. Furthermore, we find the channels produce the major error are consistent with the existence of the outliers.

We propose ASER (Activation Smoothing and Error Reconstruction), a low-rank compensation algorithm designed to enhance the efficacy of quantized models. Our method assesses the low-rank properties by whitening the relationship between quantization loss and its singular values based on Cholesky factorization. Then, ASER identifies these error-inducing outliers for smoothing. Finally, LoRA-style matrices are used to reconstruct the quantization error, with little computation overhead. The overall framework is shown in Figure 1. This lightweight yet powerful mechanism ensures robust performance of quantized models, which is orthogonal to any particular weight quantization method. Our experiments demonstrate ASER remarkably recovers the quantization error, where W4A8 per-channel quantized model performs nearly equal capability compared to the half-precision reference. This work makes the following contributions:

- We formulate the optimization objective of minimizing the discrepancy in model outputs, and analyze the characteristics of quantization errors in LLMs quantization.
- We propose a novel algorithm ASER, which includes error reconstruction using whitening SVD, and activation smoothing with outlier analysis.
- Experimental results show ASER can significantly recover the performance of quantized model in W4A8 per-channel quantization, with little computational overhead.

*These authors contributed equally.

[†]Corresponding Author.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

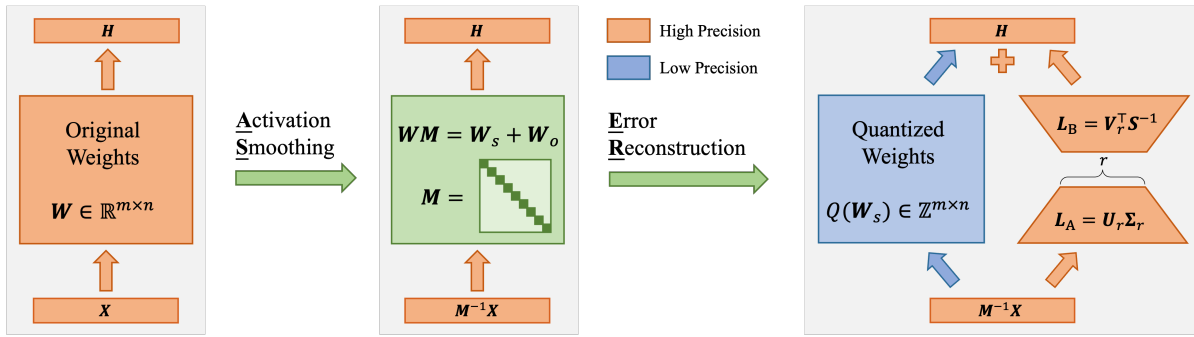


Figure 1: The framework of ASER. The LoRA-style matrices L_A, L_B are generated to reconstruct the quantization error.

Related Work

LLMs Compression and Quantization. In recent years, given the powerful emergent capabilities (Schaeffer, Miranda, and Koyejo 2024) of large language models (LLMs), there has been a growing number of model compression techniques specifically developed for Transformer-based models. Network pruning (Ma, Fang, and Wang 2023), including structured (Ma, Fang, and Wang 2023) and unstructured pruning (Kurtic et al. 2022), removes unnecessary neurons or structures to decrease model size, while distillation (Sun et al. 2020) directly uses a smaller model to learn the knowledge from the origin model. Though these conventional compression techniques have indeed yielded some benefits in LLMs, they pose challenges in deployment due to the modifications of model structure, and they often necessitate substantial consumption for retraining the models.

Post-Training Quantization (PTQ) has emerged as a foundational technology for deploying large models, overcoming these limitations and significantly enhancing inference speed. GPTQ (Frantar et al. 2022), as an extension of OBQ (Frantar and Alistarh 2022), employs second-order methods to obtain a closed-form solution, theoretically ensuring the minimization of quantization error. AWQ (Lin et al. 2024) employs a scaling factor to reduce the quantization error of salient weights. But weight-only quantization methods do not consider decreasing the bit of activations to achieve further speedup. SmoothQuant (Xiao et al. 2023) uses empirical transformations to shift the quantization difficulty from activations to weights. However, it doesn't establish the relationship between transition matrix and quantization error, and pollutes the smoothness of weights.

Low Rank Adapters of LLMs. The low-rank property is a prominent characteristic of LLMs, and researchers have attempted to structurally compress models based on this observation. Low-rank decomposition is a critical technique for compression, which has been widely utilized in LLMs (Wang et al. 2024). Linformer (Wang et al. 2020) introduces a low-rank attention mechanism, reducing the original $O(N^2)$ computational complexity. For decoder models, DeepSeek-V2 (DeepSeek-AI 2024) exploits the low-rank property of KV cache to adopt the MLA architecture, which effectively compresses the storage of the KV cache.

Another influential technique is LoRA (Hu et al. 2021), which establishes a fundamental paradigm for large model

fine-tuning. This approach of using low-rank matrices has been adopted across various domains. Recently, works like LoftQ (Li et al. 2023) and LQ-LoRA (Guo et al. 2023) employ low-rank approximation to facilitate more efficient fine-tuning of lightweight models. LoRC (Yao et al. 2024) employs them to improve the quality recovery of quantized models. L²QER (Zhang et al. 2024) further scales the quantization error by empirically designed diagonal matrix before reconstruction. We are inspired by these LoRA-style methods and propose a more compact quantization framework. ASER designs whitening SVD for error reconstruction and smoothes activations via outlier analysis.

Analysis of Quantization Error

In this section, we formulate the quantization error from a model compression perspective, and anchor minimizing the quantization loss as the objective of our problem. We first analyze the properties of the quantization loss in the context of RTN (Round-To-Nearest) quantization, and then discuss how to reconstruct it with a small number of parameters.

Minimizing the Quantization Error

Tensors in quantized models have restricted range of representations, which brings out an inevitable loss of dequantized data. In contrast to methods like LoRC that only account for the error from weight quantization, we aim to consider the integral error introduced by both weights and activations. Related research (Xiao et al. 2023) acknowledges the impact of activations on quantization, primarily due to the critical role of outliers in the activations within LLMs (Tang et al. 2024). Thus, as a typical model compression, we formulate quantization objective optimization as:

$$\arg \min_{\mathbf{W}_q} \|\mathbf{W}\mathbf{X} - \mathbf{W}_q\mathbf{X}\|_F, \quad (1)$$

where $\mathbf{W}_q = Q(\mathbf{W})$ is the quantized weights, and we use Frobenius norm to measure the magnitude of the error matrix. Here, we use $\mathbf{E}_q = \mathbf{W} - \mathbf{W}_q$ to represent the quantization error of the weight matrix. We aim to estimate it with an approximation $\tilde{\mathbf{E}}_q$, and maximize the compensation for the overall quantization error $\tilde{\mathbf{W}} = \mathbf{W}_q + \tilde{\mathbf{E}}_q$. Thus, the optimization objective can be transformed into:

$$\arg \min_{\tilde{\mathbf{E}}_q} \|(\mathbf{E}_q - \tilde{\mathbf{E}}_q)\mathbf{X}\|_F. \quad (2)$$

Characteristics of Quantization Error

We conduct empirical studies on \mathbf{E}_q to guide our reconstruction strategy for the quantization error. We take Round-To-Nearest (RTN) quantization on LLaMA3-8B (Touvron et al. 2023) as an example to obtain the layer-wise distribution.

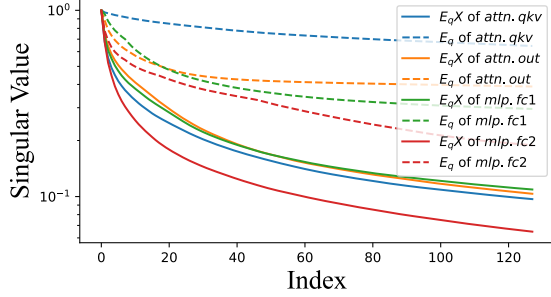


Figure 2: The largest 128 singular value distribution of quantization error in the 30th layer of LLaMA3-8B by RTN quantization. The singular values are normalized for comparison.

Activation-weight quantization error exhibits low-rank properties. As Figure 2 shows, we plot the normalized singular value distribution of \mathbf{E}_q and $\mathbf{E}_q\mathbf{X}$ over four consecutive linear layers, i.e., *qkv_proj*, *out_proj* of the multi-head self-attention (MHSA), and *fc1*, *fc2* of the feed-forward network (FFN), in a certain Transformer block. We sort them and display the top 128 largest ones out of 4096 singular values. Compared to the weight error \mathbf{E}_q , the distribution of activation-weight quantization error $\mathbf{E}_q\mathbf{X}$ features low-rank property, with a few larger values and a long tail with many smaller ones. This phenomenon is also consistent with the observation of (Yu and Wu 2023) that the features in Transformers are low-rank rather than the weights. Therefore, we can leverage this property to estimate the quantization error with low-rank approximation, as the approach of low-rank adapters (Hu et al. 2021). Especially, in consideration of $\mathbf{E}_q\mathbf{X}$, we reconstruct the data-aware integral error rather than the weights themselves like (Yao et al. 2024) does.

Low-rankness varies intra- and inter- layers. We introduce a new metric to quantify the amount of effective dimensionality in a low-rank matrices, effective rank (Roy and Vetterli 2007). Suppose the matrix \mathbf{Z} has singular values $\sigma_1, \sigma_2, \dots, \sigma_n$, its effective rank can be computed as the entropy of the normalized singular values:

$$\text{Eff_Rank}(\mathbf{Z}) = \exp\left(-\sum_{k=1}^n p_k \log p_k\right), \quad (3)$$

$$p_k = \frac{\sigma_k(\mathbf{Z})}{\sum_{i=1}^n \sigma_i(\mathbf{Z})} + \epsilon. \quad (4)$$

We plot the activation-weight quantization error $\mathbf{E}_q\mathbf{X}$ across layers in Figure 3 according to the definition of effective

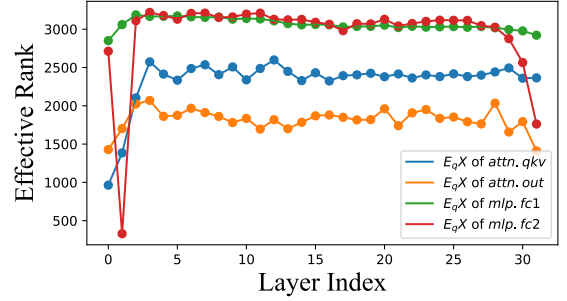


Figure 3: The effective rank of $\mathbf{E}_q\mathbf{X}$ across layers in LLaMA3-8B by RTN quantization.

rank in Eq.3. As we can see, the quantization error of linear layers in the MHSA tend to be more low-rank, whereas the counterparts in FFN have relatively higher effective dimensionality. Additionally, when examining the behavior across Transformer layers in LLaMA3-8B, we find that the quantization errors are more strongly low-rank in the layers closer to the input. This might be related to the smoother numerical distribution in the early layers of the model, which could contribute to this characteristic. Therefore, it is essential to adaptively choose appropriate ranks and the dimensions of narrow matrices to compensate for the quantization loss in each individual linear layer.

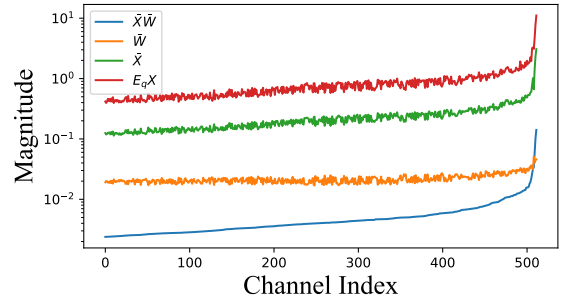


Figure 4: Magnitude of activation-weight quantization error $\mathbf{E}_q\mathbf{X}$, mean activation $\bar{\mathbf{X}}$, mean weight $\bar{\mathbf{W}}$ and their product $\bar{\mathbf{X}}\bar{\mathbf{W}}$. The channel index is sorted by $\bar{\mathbf{X}}\bar{\mathbf{W}}$. Data comes from the largest 512 channels of a linear layer in LLaMA3-8B when conducting RTN quantization.

Outliers constitute the major component of quantization error. Increasing research (Dettmers et al. 2022; Xiao et al. 2023; Ma, Fang, and Wang 2023) has recognized the significant impact of outliers on model quantization, making their handling a key consideration in quantization algorithms. In our experiments, we also observed that both activation and weight channel outliers contribute significantly to the large components in quantization error. As shown in Figure 4, when we sort the channels by $\bar{\mathbf{X}}\bar{\mathbf{W}}$, the Frobenius norm of the quantization error has also demonstrated a broadly similar trend. Notably, we find that a small fraction (less than 1% of the 4096) of outlier channels are responsible for an order

of magnitude more quantization error than the rest, and that these channels have large magnitude in both weights and activations. As a result, we attempt to implement targeted quantization and compensation strategies for them.

ASER: Activation Smoothing and Error Reconstruction

In this section, we introduce ASER, which consists of two main technologies. Error Reconstruction (ER) constructs the compensation matrices using a whitening method whose singular values are directly related to the integral error. Activation Smoothing (AS) analyzes outliers and transforms fluctuations in activation to weights, making low-rank compensation more effective.

Error Reconstruction via Whitening SVD

Firstly, we illustrate how to process quantization error using the whitening technique and establish the mapping between singular values and compression loss. To do this, we first obtain an orthogonal whitened activation $\mathbf{S}^{-1}\mathbf{X}$ where each channel is independent from each other:

$$(\mathbf{S}^{-1}\mathbf{X})(\mathbf{S}^{-1}\mathbf{X})^\top = \mathbf{S}^{-1}\mathbf{X}\mathbf{X}^\top(\mathbf{S}^{-1})^\top = \mathbf{I}, \quad (5)$$

where \mathbf{S} can be derived by Cholesky decomposition (Meyer 2000). Then we perform SVD on $\mathbf{E}_q\mathbf{S}$, which produces $\mathbf{E}_q\mathbf{S} = \mathbf{U}\Sigma\mathbf{V}^\top$, where $\mathbf{U} = [\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_n]$, $\Sigma = \text{diag}(\sigma_1, \sigma_2, \dots, \sigma_n)$ and $\mathbf{V} = [\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n]^\top$. Suppose we would approximate \mathbf{E}_q by rank r , the smallest $n-r$ singular values of Σ will be truncated to obtain the compressed $\tilde{\mathbf{E}}_q$:

$$\tilde{\mathbf{E}}_q = \mathbf{U}_r \Sigma_r \mathbf{V}_r^\top \mathbf{S}^{-1}. \quad (6)$$

Next, we provide a theoretical justification establishing the connection between singular values and error compensation, showing that the top singular values indeed contribute the most to mitigating the overall error.

For an arbitrary matrix $\mathbf{A} \in \mathbb{R}^{m \times m}$, its Frobenius norm can be deduced into the square root of the trace of its gram matrix (Horn and Johnson 1991):

$$\|\mathbf{A}\|_F = \left(\sum_{j=1}^n \sum_{i=1}^m |a_{ij}|^2 \right)^{\frac{1}{2}} = \left[\text{Tr}(\mathbf{A}^\top \mathbf{A}) \right]^{\frac{1}{2}}, \quad (7)$$

based on which we can derive the approximation loss L_i of $\tilde{\mathbf{E}}_q\mathbf{X}$ when truncating the i^{th} singular value:

$$\begin{aligned} L_i &= \|(\mathbf{E}_q - \tilde{\mathbf{E}}_q)\mathbf{X}\|_F = \|\sigma_i \mathbf{u}_i \mathbf{v}_i^\top \mathbf{S}^{-1} \mathbf{X}\|_F \\ &\stackrel{(i)}{=} \sigma_i \text{Tr}(\mathbf{u}_i \mathbf{v}_i^\top \mathbf{S}^{-1} \mathbf{X} \mathbf{X}^\top (\mathbf{S}^{-1})^\top \mathbf{v}_i \mathbf{u}_i^\top)^{\frac{1}{2}} \\ &\stackrel{(ii)}{=} \sigma_i \text{Tr}(\mathbf{u}_i \mathbf{v}_i^\top \mathbf{v}_i \mathbf{u}_i^\top)^{\frac{1}{2}} \stackrel{(iii)}{=} \sigma_i, \end{aligned} \quad (8)$$

where (i) refers to Eq. 7, (ii) refers to Eq. 6 and (iii) equals since $\mathbf{v}_i^\top \mathbf{v}_i = \mathbf{u}_i^\top \mathbf{u}_i = I$; $\mathbf{v}_i^\top \mathbf{v}_j = \mathbf{u}_i^\top \mathbf{u}_j = 0, \forall i \neq j$; $\text{Tr}(\mathbf{v}_i \mathbf{v}_i^\top) = \text{Tr}(\mathbf{u}_i \mathbf{u}_i^\top) = 1$.

Therefore, $\tilde{\mathbf{E}}_q$ can be replaced with a pair of LoRA-style matrices $\mathbf{L}_A = \mathbf{U}_r \Sigma_r, \mathbf{L}_B = \mathbf{V}_r^\top \mathbf{S}^{-1}$ to compensate for the

quantization error of the model. In addition to setting the same rank r for all layers, we use the ratio of the cumulative sum of sorted singular values to the total sum of singular values to set a threshold α whose range is (0, 1).

$$\max_r \frac{\sum_{i=1}^r \sigma_i}{\sum_{i=1}^n \sigma_i} < \alpha. \quad (9)$$

Activation Smoothing via Outlier Analysis

Based on the empirical observations presented in the analysis section, we find that the outliers play a critical role in error reconstruction. Therefore, we devise a heuristic method to process these outliers. The activation outliers are extracted and migrated to the weights, after which the outlier parts of weights can be split and compensated along with the quantization error reconstruction.

Firstly, inspired by SmoothQuant (Xiao et al. 2023), we use a scaling matrix M to migrate the quantization difficulty of the activation to weight:

$$\mathbf{W}\mathbf{X} = \mathbf{W}\mathbf{M} \cdot \mathbf{M}^{-1}\mathbf{X}, \quad (10)$$

where $\mathbf{M} = \text{diag}(\mathbf{m}_1, \mathbf{m}_2, \dots, \mathbf{m}_n)$ is a specially designed matrix. Let $\bar{\mathbf{X}}$ and $\bar{\mathbf{W}}$ denote the absolute mean activation and weight value of each channel respectively. Suppose f is a pre-defined hyper-parameter, we use $I_f = \{i_1, i_2, i_3, \dots, i_f\}$ to represent the index set of $\bar{\mathbf{X}} \odot \bar{\mathbf{W}}$ outliers. \mathbf{M} can be obtained through the following formula:

$$\mathbf{m}_i = \begin{cases} \bar{\mathbf{X}}_i / \bar{\mathbf{X}}_{\min} & i \in I_f \\ 1 & \text{else,} \end{cases} \quad (11)$$

where $\bar{\mathbf{X}}_{\min}$ is the minimum value in $\bar{\mathbf{X}}_{I_f}$. This makes the activation values easier to quantize and meanwhile allows us to use whitening SVD to uniformly treat and compensate for integral quantization error in the weight matrix:

$$\begin{aligned} \mathbf{W}\mathbf{M} \cdot \mathbf{M}^{-1}\mathbf{X} &= (\mathbf{W}_s + \mathbf{W}_o)\mathbf{M}^{-1}\mathbf{X} \\ &= (Q(\mathbf{W}_s) + (\mathbf{E}_q + \mathbf{W}_o))\mathbf{M}^{-1}\mathbf{X}. \end{aligned} \quad (12)$$

The outlier part \mathbf{W}_o will be extracted from the adjusted weight $\mathbf{W}\mathbf{M}$ based on the outlier index set I_f , and it will not be quantized. Together with the quantization errors from the rest of the weights, they are compensated by low-rank approximation matrices to minimize the loss incurred during the quantization process. Thus, we morph the objective of the whitening SVD into $(\mathbf{E}_q + \mathbf{W}_o)$, and similar narrow matrices are used to approximate this error:

$$\mathbf{L}_A \mathbf{L}_B \sim (\mathbf{E}_q + \mathbf{W}_o)\mathbf{S}. \quad (13)$$

In this way, the inference pass is replaced by $\mathbf{W}_q\mathbf{X} + \mathbf{L}_A \mathbf{L}_B \mathbf{X}$ for error reconstruction. Note that this technique is orthogonal to the weight-only quantization, which means ASER can be conducted based on the existing weight-only quantization methods, not limited to RTN. To sum up, the pseudo-code of our algorithm is as follows:

Algorithm 1: Activation Smoothing and Error Reconstruction for Post-Training Quantization

Input: Original weight \mathbf{W} , Calibration data \mathcal{D}
Parameter: Rank threshold α , Outlier threshold f
Output: Quantized Weight \mathbf{W}_q , Reconstruction matrices $\mathbf{L}_A, \mathbf{L}_B$ over L layers

- 1: **for** Data d in \mathcal{D} **do**
- 2: Inference data d and store activations \mathbf{X}
- 3: **end for**
- 4: **for** Layer \mathbf{W}^l in \mathbf{W} **do**
- 5: Initialize $\mathbf{m} = [1, 1, \dots, 1]^n$ and $\mathbf{M} = \text{diag}(\mathbf{m})$
- 6: **if** with Activation Smoothing (A.S.) **then**
- 7: Compute $\mathbf{m}_i = \bar{\mathbf{X}}_i^l / \bar{\mathbf{X}}_{\min}^l$, for $i \in I_f$
- 8: Split $\mathbf{W}^l \mathbf{M} = \mathbf{W}_s^l + \mathbf{W}_o^l$
- 9: Compute $\mathbf{E}_q^l = \mathbf{W}^l - Q(\mathbf{W}_s^l) + \mathbf{W}_o^l$
- 10: **else**
- 11: Compute $\mathbf{E}_q^l = \mathbf{W}^l - Q(\mathbf{W}^l)$
- 12: **end if**
- 13: Compute \mathbf{S} by Cholesky decomposition on $(\mathbf{M}^{-1}\mathbf{X})(\mathbf{M}^{-1}\mathbf{X})^\top$: $(\mathbf{S}^{-1}\mathbf{M}^{-1}\mathbf{X})(\mathbf{S}^{-1}\mathbf{M}^{-1}\mathbf{X})^\top = \mathbf{I}$
- 14: SVD decomposition: $\mathbf{U}\mathbf{\Sigma}\mathbf{V}^\top = \mathbf{E}_q^l \mathbf{S}$
- 15: Compute r according to Eq. 9
- 16: Compute $\mathbf{L}_A^l = \mathbf{U}_r \mathbf{\Sigma}_r, \mathbf{L}_B^l = \mathbf{V}_r^\top \mathbf{S}^{-1}$
- 17: **end for**
- 18: **return** $\{\mathbf{W}_q^l, \mathbf{L}_A^l, \mathbf{L}_B^l\}_{l=1}^L$

Experiments

This section presents experimental results by different PTQ methods on various model architectures. We further conduct visualization analysis to validate the effectiveness of the method, followed by an overhead analysis.

Experimental Setup

Models We consider quantization of typical open source large language models involving LLaMA3-8B (Touvron et al. 2023), Qwen1.5-7B, Qwen-72B (Bai et al. 2023). We obtain the pre-trained model from the official repository¹².

Baselines We consider several popular PTQ algorithms for comparison. The baseline approaches include act-and-weight quantization approaches LLM.int() (Dettmers et al. 2022), SmoothQuant (Xiao et al. 2023) and SmoothQuant+ (Pan et al. 2023). Furthermore, we compared the quantization algorithm LoRC (Yao et al. 2024), L²QER (Zhang et al. 2024) using additional parameter. In our results, we use w/ and w/o A.S. to indicate whether the activation smoothing has been employed.

Evaluation Metrics Firstly, we evaluate quantized model by perplexity (PPL) on language modeling datasets Wikitext2 (Merity et al. 2016), PTB (Marcus et al. 1994) and a portion of the C4 dataset (Raffel et al. 2020). Then we conduct experiments on zero-shot evaluation tasks including ARC-e, ARC-c (Clark et al. 2018), MMLU (Hendrycks et al. 2020), HellaSwag (Zellers et al. 2019), PIQA (Bisk et al.

2020), GSM8K (Cobbe et al. 2021) and Human-Eval (Chen et al. 2021). We experiment on the open source evaluation framework OpenCompass (OpenCompass 2023).

Experimental Details To ensure a fair comparison in our assessment, we standardize the usage of the calibration dataset for all quantization methods. All quantization is conducted by per-channel for weights and per-token for activations. We consistently employ a sample size of 128 and maintain a uniform token sequence length of 2048 for calibration throughout the evaluation process. All experiments are conducted on 4× NVIDIA A100 GPUs.

Main Results

We demonstrate the main results of LLaMA3-8B and Qwen1.5-7B evaluation in Tables 1 and 2 respectively. The methods are divided into two parts, W4A8 and W4A6 quantization setups. We both test the perplexity and accuracy, and the best results are marked in bold. To ensure a fair comparison, we set the rank of compensation parameters in each layer of ASER, LoRC, and L²QER to 64. The outlier threshold f is set by 32, which is equal in experimental setup.

As shown in the table 1, ASER achieves the best perplexity on language modeling tasks, while simultaneously performs competitive accuracy on language understanding. In both settings, ASER achieves perplexity improvement compared to the fp16 baseline ranging from 1.20 to 3.88. For the accuracy evaluation, particularly in the W4A6 per-channel setting, existing algorithms without error compensation, e.g., SmoothQuant and its variant, show significant performance degradation, while ASER outperforms its counterparts L²QER by 4.87%, which attributes to the significant role of outlier extraction and activation smoothing.

In Table 2, for the W4A8 quantization, ASER shows competitive performance compared to the state-of-the-art method, with 3.12% close to the benchmark accuracy of the fp16 baseline averagely. Thanks to the activation smoothing, ASER achieves 5.09% higher accuracy compared to the one with only error reconstruction. For W4A6 quantization setup on language modeling datasets, as we can see, many existing methods fail to keep the performance without group-wise scaling. However, the quantized model using ASER achieves the perplexity that is 3.08 ~ 4.19 close to the fp16 baseline. Similarly, it also delivers the best accuracy on the commonsense evaluation dataset. Particularly, compared to L²QER, ASER performs better accuracy by 6.66%, which also far exceeds other act-and-weight quantization baselines.

Ablation Study and Analysis

Results of Larger Model Next, we validate the effectiveness of our method on the models with larger-scale. The model involved is Qwen-72B in Table 3. We mainly evaluate the setup of W4A8 per-channel quantization on commonsense benchmarks. As we can see, relying on our low-rank quantization error reconstruction algorithm, with further activation smoothing, ASER achieves an accuracy of 0.33% ~ 4.27%, which is close to the fp16 baseline. This demonstrates the scalability of our approach and its adaptability to larger models. Experimental results of more quantization setups and models and can be found in Appendix.

¹<https://github.com/meta-llama/llama>

²<https://github.com/QwenLM/Qwen>

Method	#W	#A	Perplexity(\downarrow)			Accuracy(\uparrow)					
			WikiText2	C4	PTB	ARC-e	ARC-c	MMLU	Hella	PIQA	Avg.
LLaMA3-8B	16	16	6.14	10.75	9.91	89.77	79.66	37.85	48.47	68.55	64.86
LLM.int4()	4	8	10.57	18.43	15.52	77.07	59.32	13.48	28.83	54.68	46.68
SmoothQuant	4	8	11.23	18.87	16.78	60.32	44.75	14.99	26.93	37.65	36.93
SmoothQuant+	4	8	10.22	17.39	15.34	72.31	55.93	14.25	26.44	55.17	44.82
LoRC	4	8	8.62	15.09	12.99	83.07	65.08	10.38	29.79	60.99	49.86
L ² QER	4	8	7.72	13.51	11.84	86.07	69.15	26.83	30.42	61.97	54.89
ASER (w/o A.S.)	4	8	7.61	13.24	11.65	86.24	72.20	27.90	31.08	59.47	55.38
ASER (w/ A.S.)	4	8	7.43	12.94	11.36	86.24	72.54	30.73	29.62	60.50	55.93
LLM.int4()	4	6	14.21	24.52	22.89	51.85	40.68	8.30	24.54	49.18	34.91
SmoothQuant	4	6	15.74	25.29	24.42	39.51	30.17	7.58	23.17	41.73	28.43
SmoothQuant+	4	6	29.22	39.10	42.46	30.16	32.54	15.51	19.66	41.64	27.90
LoRC	4	6	11.03	19.28	16.99	61.20	42.04	3.43	25.94	54.35	37.39
L ² QER	4	6	9.44	16.43	14.12	70.37	56.27	10.98	26.48	54.30	43.68
ASER (w/o A.S.)	4	6	9.02	15.55	13.46	73.37	52.20	7.78	26.99	55.55	43.18
ASER (w/ A.S.)	4	6	8.41	14.63	12.58	78.66	61.69	17.32	27.07	58.00	48.55

Table 1: Evaluation results of post-training quantization on LLaMA3-8B. Best results are marked in bold.

Method	#W	#A	Perplexity(\downarrow)			Accuracy(\uparrow)					
			WikiText2	C4	PTB	ARC-e	ARC-c	MMLU	Hella	PIQA	Avg.
Qwen1.5-7B	16	16	7.95	13.57	11.94	87.48	76.61	49.78	61.49	66.76	68.42
LLM.int4()	4	8	23.73	38.14	25.28	17.46	14.58	13.27	0.18	12.08	11.51
SmoothQuant	4	8	15.73	28.18	21.24	33.51	26.44	24.00	1.88	31.94	21.46
SmoothQuant+	4	8	46.27	51.28	40.81	28.04	22.03	9.16	1.17	29.60	18.00
LoRC	4	8	16.78	27.13	19.93	76.19	66.44	21.56	43.77	45.43	50.68
L ² QER	4	8	9.20	15.63	13.81	85.36	71.53	46.39	29.59	64.64	59.50
ASER (w/o A.S.)	4	8	9.19	15.59	13.69	81.83	68.81	43.89	45.44	61.10	60.21
ASER (w/ A.S.)	4	8	8.72	14.84	13.10	84.66	72.54	49.55	52.18	67.57	65.30
LLM.int4()	4	6	34.32	56.00	39.67	50.09	39.66	20.89	16.11	35.69	32.49
SmoothQuant	4	6	26.81	42.67	35.62	41.98	35.59	16.08	7.31	32.59	26.71
SmoothQuant+	4	6	131.45	176.46	130.24	35.98	30.51	17.36	18.77	33.35	27.19
LoRC	4	6	23.03	39.63	29.68	45.68	37.63	8.90	26.29	30.85	29.87
L ² QER	4	6	16.51	27.01	22.46	63.14	47.80	21.85	30.24	51.52	42.91
ASER (w/o A.S.)	4	6	15.86	26.38	21.25	67.37	54.92	40.34	30.56	54.57	49.52
ASER (w/ A.S.)	4	6	11.03	17.57	16.13	75.31	62.37	32.07	34.49	43.63	49.57

Table 2: Evaluation results of post-training quantization on Qwen1.5-7B. Best results are marked in bold.

Method	Accuracy(\uparrow)			
	ARC-e	ARC-c	GSM8K	HEval
Qwen-72B	94.36	88.47	70.13	50.61
LLM.int4()	92.77	84.75	56.33	39.02
SmoothQuant	72.13	65.08	56.71	38.41
SmoothQuant+	86.07	77.63	60.58	37.20
LoRC	93.84	88.14	64.06	45.12
L ² QER	92.95	85.46	66.19	42.07
ASER (w/o AS)	94.00	88.14	64.14	44.51
ASER (w/ AS)	92.59	85.76	67.32	46.34

Table 3: Evaluation results on Qwen-72B. All models are quantized by W4A8 per-channel setup. The baseline model is parameterized by fp16. HEval represents the Human-Eval benchmark. Best results are marked in bold.

Effect of Activation Smoothing As previously mentioned, ASER employs several techniques, including activation smoothing, migration of outliers, and low-rank estimation to compensate for integral quantization errors. These techniques make activation quantization more favorable. The results of the main experiments demonstrate the improvement brought by activation smoothing particularly in activation quantization setup. So we further conduct some experiments to study this effect. As Figure 5 shows, the existing quantization algorithms reach intolerable perplexity with low bit activation. LLM.int4(), SmoothQuant and its variant, cannot withstand the performance degradation caused by low activation bit-width. Though LoRC and L²QER do compensation for the quantized weight, they are not able to process activation outliers. However, ASER keeps remarkably better performance even in 4-bit activation quantization.

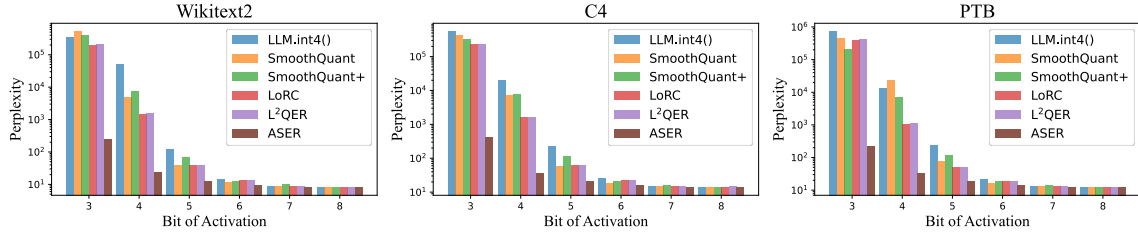


Figure 5: Perplexity of quantized Qwen1.5-7B with int8 weight and different quantization bit of activation, i.e., W8Ax.

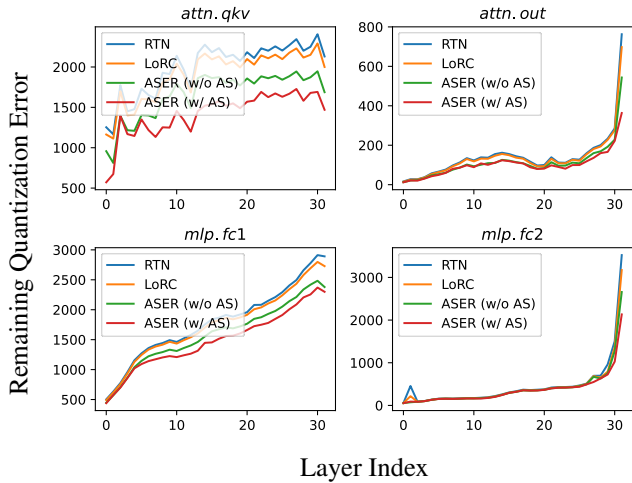


Figure 6: The remaining quantization error across layers of LLaMA3-8B using different methods of W4A6 quantization. The error is $\|\mathbf{W}\mathbf{X} - \mathbf{W}_q\mathbf{X}_q\|_F$ for the RTN, while the other methods represent the error norm after reconstruction.

Effect of Error Compensation We visualize the compensation effect of different methods in Fig. 6, where the remaining quantization errors of four consecutive layers are plotted over transformer layers in LLaMA3-8B. The RTN represents the baseline quantization error without any compensation. ASER can effectively reconstruct the quantization error, compared to LoRC. Furthermore, the application of activation smoothing enhances the error reconstruction capability due to the extraction and processing of outliers.

Rank Selection In the ASER framework, the choice of rank r , or its threshold α , influences the amount of quantization error compensation and the additional number of parameters introduced. Therefore, we conduct an experimental analysis on how the selection of rank affects downstream tasks. For Qwen1.5-7B, the performance of downstream tasks positively correlates with the selection of α , whereas this might not be the case for other models. As shown in Table 4, we find that using a larger rank is not necessarily better for all tasks in Qwen-7B. When we set smaller rank, we might achieve a better balance between performance of W4A8-quantized model and the extra parameter count.

α	\bar{r}	ARC-e	Hella	PIQA	+ FLOPs
0.1	59.93	83.25	68.39	72.52	2.93%
0.075	49.60	84.48	67.63	72.74	2.42%
0.05	30.58	82.36	67.33	72.80	1.49%
0.03	14.00	80.78	67.41	72.96	0.68%
0.015	5.41	82.89	67.01	71.98	0.26%

Table 4: Zero-shot accuracy evaluation results and overhead of Qwen-7B using ASER with different rank configuration.

Overhead Analysis

The proposed method can introduce additional computations and memory overheads due to the involvement of extra LoRA-style matrices $\mathbf{L}_A, \mathbf{L}_B$. Let s denote the sequence length, d the hidden dimension, r ($\ll d$) the rank. Table 5 analyzes the complexity of a single layer, where the computation represents floating-point operations (FLOPs), and the memory is measured by the number of parameters.

Complexity	Origin	ASER
Computation	sd^2	$sd^2 + 2srd$
Memory	d^2	$d^2 + 2rd$

Table 5: Complexity Analysis of ASER quantized model.

Taking Qwen-7B as an example, we demonstrate the tradeoff of accuracy and the computational overhead among the listed configurations of Table 4. The choice can be deliberated based on specific model and tasks. The additional parameter introduces theoretical computational overhead of only 0.26%, which is minor to the main network inference. We leave it for future work to adaptively choose the best α .

Conclusion

In this work, from the perspective of model compression, we formulate and analyze the quantization error. Our study reveals that the error exhibits low-rank property, whose major components originate from the combined effect of outliers in both activations and weights. In response to these findings, we propose ASER, which utilizes whitening SVD to compensate for the integral error (Error Reconstruction), followed by analyzing and extracting activation outliers (Activation Smoothing). Experimental results demonstrate that ASER outperforms existing quantization algorithms on typical large language models in various quantization setups, with negligible theoretical overhead.

References

- Bai, J.; Bai, S.; Chu, Y.; Cui, Z.; Dang, K.; Deng, X.; Fan, Y.; Ge, W.; Han, Y.; Huang, F.; et al. 2023. Qwen technical report. *arXiv preprint arXiv:2309.16609*.
- Bisk, Y.; Zellers, R.; Gao, J.; Choi, Y.; et al. 2020. Piqa: Reasoning about physical commonsense in natural language. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, 7432–7439.
- Chen, M.; Tworek, J.; Jun, H.; Yuan, Q.; de Oliveira Pinto, H. P.; Kaplan, J.; Edwards, H.; Burda, Y.; Joseph, N.; Brockman, G.; Ray, A.; Puri, R.; Krueger, G.; Petrov, M.; Khlaaf, H.; Sastry, G.; Mishkin, P.; Chan, B.; Gray, S.; Ryder, N.; Pavlov, M.; Power, A.; Kaiser, L.; Bavarian, M.; Winter, C.; Tillet, P.; Such, F. P.; Cummings, D.; Plappert, M.; Chantzis, F.; Barnes, E.; Herbert-Voss, A.; Guss, W. H.; Nichol, A.; Paino, A.; Tezak, N.; Tang, J.; Babuschkin, I.; Balaji, S.; Jain, S.; Saunders, W.; Hesse, C.; Carr, A. N.; Leike, J.; Achiam, J.; Misra, V.; Morikawa, E.; Radford, A.; Knight, M.; Brundage, M.; Murati, M.; Mayer, K.; Welinder, P.; McGrew, B.; Amodei, D.; McCandlish, S.; Sutskever, I.; and Zaremba, W. 2021. Evaluating Large Language Models Trained on Code.
- Clark, P.; Cowhey, I.; Etzioni, O.; Khot, T.; Sabharwal, A.; Schoenick, C.; and Tafford, O. 2018. Think you have solved question answering? try arc, the ai2 reasoning challenge. *arXiv preprint arXiv:1803.05457*.
- Cobbe, K.; Kosaraju, V.; Bavarian, M.; Chen, M.; Jun, H.; Kaiser, L.; Plappert, M.; Tworek, J.; Hilton, J.; Nakano, R.; Hesse, C.; and Schulman, J. 2021. Training Verifiers to Solve Math Word Problems. *arXiv preprint arXiv:2110.14168*.
- DeepSeek-AI. 2024. DeepSeek-V2: A Strong, Economical, and Efficient Mixture-of-Experts Language Model. *arXiv:2405.04434*.
- Dettmers, T.; Lewis, M.; Belkada, Y.; and Zettlemoyer, L. 2022. LLM.int8(): 8-bit Matrix Multiplication for Transformers at Scale. *arXiv:2208.07339*.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Frantar, E.; and Alistarh, D. 2022. Optimal brain compression: A framework for accurate post-training quantization and pruning. *Advances in Neural Information Processing Systems*, 35: 4475–4488.
- Frantar, E.; Ashkboos, S.; Hoefler, T.; and Alistarh, D. 2022. Gptq: Accurate post-training quantization for generative pre-trained transformers. *arXiv preprint arXiv:2210.17323*.
- Guo, H.; Greengard, P.; Xing, E. P.; and Kim, Y. 2023. Lq-lora: Low-rank plus quantized matrix decomposition for efficient language model finetuning. *arXiv preprint arXiv:2311.12023*.
- Hendrycks, D.; Burns, C.; Basart, S.; Zou, A.; Mazeika, M.; Song, D.; and Steinhardt, J. 2020. Measuring massive multitask language understanding. *arXiv preprint arXiv:2009.03300*.
- Horn, R. A.; and Johnson, C. R. 1991. *Topics in Matrix Analysis*. Cambridge University Press.
- Hu, E. J.; Shen, Y.; Wallis, P.; Allen-Zhu, Z.; Li, Y.; Wang, S.; Wang, L.; and Chen, W. 2021. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*.
- Kurtic, E.; Campos, D.; Nguyen, T.; Frantar, E.; Kurtz, M.; Fineran, B.; Goin, M.; and Alistarh, D. 2022. The optimal bert surgeon: Scalable and accurate second-order pruning for large language models. *arXiv preprint arXiv:2203.07259*.
- Li, Y.; Yu, Y.; Liang, C.; He, P.; Karampatziakis, N.; Chen, W.; and Zhao, T. 2023. Loftq: Lora-fine-tuning-aware quantization for large language models. *arXiv preprint arXiv:2310.08659*.
- Lin, J.; Tang, J.; Tang, H.; Yang, S.; Chen, W.-M.; Wang, W.-C.; Xiao, G.; Dang, X.; Gan, C.; and Han, S. 2024. AWQ: Activation-aware Weight Quantization for On-Device LLM Compression and Acceleration. *Proceedings of Machine Learning and Systems*, 6: 87–100.
- Liu, Z.; Zhao, C.; Fedorov, I.; Soran, B.; Choudhary, D.; Krishnamoorthi, R.; Chandra, V.; Tian, Y.; and Blanckevoort, T. 2024. SpinQuant-LLM quantization with learned rotations. *arXiv preprint arXiv:2405.16406*.
- Llama Team, A. . M. 2024. The Llama 3 Herd of Models. *arXiv:2407.21783*.
- Ma, X.; Fang, G.; and Wang, X. 2023. Llm-pruner: On the structural pruning of large language models. *Advances in neural information processing systems*, 36: 21702–21720.
- Marcus, M.; Kim, G.; Marcinkiewicz, M. A.; MacIntyre, R.; Bies, A.; Ferguson, M.; Katz, K.; and Schasberger, B. 1994. The penn treebank: Annotating predicate argument structure. In *Human Language Technology: Proceedings of a Workshop held at Plainsboro, New Jersey, March 8-11, 1994*.
- Merity, S.; Xiong, C.; Bradbury, J.; and Socher, R. 2016. Pointer sentinel mixture models. *arXiv preprint arXiv:1609.07843*.
- Meyer, C. D. 2000. Matrix Analysis and Applied Linear Algebra Book and Solutions Manual. *Society for Industrial and Applied Mathematics*.
- OpenCompass. 2023. OpenCompass: A Universal Evaluation Platform for Foundation Models. <https://github.com/open-compass/opencompass>.
- Pan, J.; Wang, C.; Zheng, K.; Li, Y.; Wang, Z.; and Feng, B. 2023. SmoothQuant+: Accurate and Efficient 4-bit Post-Training WeightQuantization for LLM. *arXiv:2312.03788*.
- Raffel, C.; Shazeer, N.; Roberts, A.; Lee, K.; Narang, S.; Matena, M.; Zhou, Y.; Li, W.; and Liu, P. J. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of machine learning research*, 21(140): 1–67.
- Roy, O.; and Vetterli, M. 2007. The effective rank: A measure of effective dimensionality. In *2007 15th European signal processing conference*, 606–610. IEEE.
- Schaeffer, R.; Miranda, B.; and Koyejo, S. 2024. Are emergent abilities of large language models a mirage? *Advances in Neural Information Processing Systems*, 36.

Sun, S.; Gan, Z.; Cheng, Y.; Fang, Y.; Wang, S.; and Liu, J. 2020. Contrastive distillation on intermediate representations for language model compression. *arXiv preprint arXiv:2009.14167*.

Tang, H.; Sun, Y.; Wu, D.; Liu, K.; Zhu, J.; and Kang, Z. 2024. EasyQuant: An Efficient Data-free Quantization Algorithm for LLMs. *arXiv:2403.02775*.

Touvron, H.; Lavril, T.; Izacard, G.; Martinet, X.; Lachaux, M.-A.; Lacroix, T.; Rozière, B.; Goyal, N.; Hambro, E.; Azhar, F.; et al. 2023. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*.

Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.

Wang, S.; Li, B. Z.; Khabsa, M.; Fang, H.; and Ma, H. 2020. Linformer: Self-attention with linear complexity. *arXiv preprint arXiv:2006.04768*.

Wang, X.; Zheng, Y.; Wan, Z.; and Zhang, M. 2024. Svd-llm: Truncation-aware singular value decomposition for large language model compression. *arXiv preprint arXiv:2403.07378*.

Xiao, G.; Lin, J.; Seznec, M.; Wu, H.; Demouth, J.; and Han, S. 2023. Smoothquant: Accurate and efficient post-training quantization for large language models. In *International Conference on Machine Learning*, 38087–38099. PMLR.

Yang, A.; Yang, B.; Hui, B.; Zheng, B.; Yu, B.; Zhou, C.; Li, C.; Li, C.; Liu, D.; Huang, F.; Dong, G.; Wei, H.; Lin, H.; Tang, J.; Wang, J.; Yang, J.; Tu, J.; Zhang, J.; Ma, J.; Yang, J.; Xu, J.; Zhou, J.; Bai, J.; He, J.; Lin, J.; Dang, K.; Lu, K.; Chen, K.; Yang, K.; Li, M.; Xue, M.; Ni, N.; Zhang, P.; Wang, P.; Peng, R.; Men, R.; Gao, R.; Lin, R.; Wang, S.; Bai, S.; Tan, S.; Zhu, T.; Li, T.; Liu, T.; Ge, W.; Deng, X.; Zhou, X.; Ren, X.; Zhang, X.; Wei, X.; Ren, X.; Liu, X.; Fan, Y.; Yao, Y.; Zhang, Y.; Wan, Y.; Chu, Y.; Liu, Y.; Cui, Z.; Zhang, Z.; Guo, Z.; and Fan, Z. 2024. Qwen2 Technical Report. *arXiv:2407.10671*.

Yao, Z.; Wu, X.; Li, C.; Youn, S.; and He, Y. 2024. Exploring post-training quantization in llms from comprehensive study to low rank compensation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 19377–19385.

Yu, H.; and Wu, J. 2023. Compressing transformers: features are low-rank, but weights are not! In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 11007–11015.

Zellers, R.; Holtzman, A.; Bisk, Y.; Farhadi, A.; and Choi, Y. 2019. Hellaswag: Can a machine really finish your sentence? *arXiv preprint arXiv:1905.07830*.

Zhang, C.; Cheng, J.; Constantinides, G. A.; and Zhao, Y. 2024. LQER: Low-Rank Quantization Error Reconstruction for LLMs. *arXiv preprint arXiv:2402.02446*.