# SERQ: SALIENCY-AWARE LOW-RANK ERROR RECONSTRUCTION FOR LLM QUANTIZATION

**Anonymous authors** 

000

001

002 003 004

010 011

012

013

014

016

018

019

021

024

025

026

027

028

029

031

033 034

037

038

040

041

043

044

046

047

048

051

052

Paper under double-blind review

#### **ABSTRACT**

Post-training quantization (PTQ) has emerged as a prevailing technique for deploying large language models (LLMs) efficiently in terms of both memory and computation, across edge devices and server platforms. Existing PTQ methods primarily aim to reduce precision in weights and activations by mitigating quantization errors caused by channel-wise outlier activations (e.g., pre-quantization scaling, online transformations, or low-rank error reconstruction). Among these approaches, error reconstruction with low-rank adaptation (LoRA) has proven particularly effective, as it introduces a lightweight auxiliary computation path without requiring heavy optimization or additional online layers. However, prior studies reveal severe accuracy degradation under W4A4 settings, and conventional low-rank adaptations rely on two sequential factors, necessitating intermediate quantization during inference and thereby limiting low-precision efficiency. In this work, we propose SERQ, a saliency-aware error reconstruction method for lowbit LLM inference that employs a single low-rank compensation matrix. SERQ preserves efficient 4-bit matrix multiplication in linear layers by jointly mitigating quantization errors arising from both activation and weight saliency through three stages: (1) static activation flattening, (2) saliency-aware error reconstruction, and (3) offline weight permutation. The method incurs additional computation only for low-rank error reconstruction via a single decomposition, while all other operations are performed offline, thereby keeping latency overhead minimal. Empirically, SERQ outperforms prior error reconstruction methods under both W4A8 and W4A4 settings, and achieves higher accuracy than state-of-the-art rotationbased W4A4 approaches, while substantially reducing calibration complexity.

## 1 Introduction

The demand for efficient deployment of large language models (LLMs) has been rapidly increasing across both server and edge platforms. Quantization has emerged as one of the most effective approaches to reduce the substantial memory and computational costs associated with LLM inference. In particular, post-training quantization (PTQ) techniques (Nagel et al. (2021)) enable low-precision representations of weights and activations involved in large-scale computations, thereby avoiding expensive fine-tuning while maintaining competitive performance.

A central challenge in minimizing quantization errors for LLMs lies in addressing outlier activations across channels. To alleviate this issue, several distribution-flattening approaches have been proposed, including pre-quantization scaling methods (Xiao et al. (2024); Shao et al. (2024)) and online transformation-based techniques that leverage random Hadamard or learned transformations (Ashkboos et al. (2024); Liu et al. (2025)). While recent rotation transformation methods have demonstrated effectiveness in enabling 4-bit integer (INT4) quantization, they typically rely on computationally expensive calibration procedures or suffer from performance variability induced by random matrices, thereby limiting their practicality in general deployment.

An alternative strategy for mitigating activation outliers is matrix decomposition. Recent advances have introduced low-rank error reconstruction methods that integrate quantization with low-rank adaptation (LoRA) (Dettmers et al. (2023); Saha et al. (2024); Zhang et al. (2024); Zhao et al. (2024)). These approaches leverage low-rank decompositions in matrix multiplication to reduce quantization error by compensating for it through separate low-rank factors. For example, L<sup>2</sup>QER

(Zhang et al. (2024)) introduces a fully quantized path for low-rank error reconstruction, yielding near-zero accuracy loss under the 4-bit weight, 8-bit activation (W4A8) configuration. Despite their superior adaptability, these methods have not yet achieved W4A4 quantization without noticeable performance degradation. Furthermore, they remain unsuitable for fully low-precision execution, as decomposed matrices are multiplied sequentially, producing intermediate values, requiring an additional on-the-fly quantization process.

In this work, we propose SERQ, a saliency-aware error reconstruction method that enables low-precision LLM inference (e.g., W4A4, W4A8) using a single low-rank decomposition. Unlike standard low-rank approximations that rely on two low-rank factors, our method unifies error correction into a single matrix by jointly addressing activation and weight saliency. This design avoids the overhead of an additional sequential low-rank branch during inference, while effectively mitigating quantization errors arising from both activation outliers and salient weights. As a result, SERQ achieves more efficient low-precision inference while maintaining high accuracy, outperforming prior quantization approaches in the challenging W4A4 precision setting.

To the best of our knowledge, this is the first work to realize 4-bit matrix multiplication in linear layers by employing low-rank error reconstruction, a method recognized for its adaptability and minimal calibration overhead. Following this principle, SERQ introduces no additional layers for online processing and avoids costly calibration procedures such as hyperparameter search, or other compute-intensive training operations. Our contributions are summarized as follows.

- We propose SERQ, a novel W4A4 quantization scheme for LLMs that employs a single saliency-guided low-rank matrix for accurate error reconstruction. Our method operates in three steps: static activation flattening, saliency-aware error reconstruction, and offline weight permutation.
- SERQ enables 4-bit matrix multiplication (e.g. INT4, MXFP4) in linear layers, thereby minimizing inference overhead in low-rank computation. Moreover, the proposed flattening and permutation schemes are merged into weight parameters and preprocessed in adjacent layers, allowing them to be fully managed offline with no additional latency.
- We validate our scheme across various LLMs with comprehensive evaluations. Compared to prior LoRA-based methods, our approach achieves superior performance in both W4A8 and W4A4 configurations while using only a single low-rank matrix. Furthermore, we compare against state-of-the-art rotation-based W4A4 quantization approaches, demonstrating superior accuracy while significantly reducing calibration complexity.

## 2 BACKGROUND AND MOTIVATION

## 2.1 LLM QUANTIZATION

Quantization maps high-precision values to low-precision representations, improving both memory and compute efficiency. The basic integer quantization with max-scaling can be expressed as:

$$X_q = \operatorname{clip}(\lceil X/s \rfloor), \quad s = \max(|X|)/(2^{n-1} - 1), \quad \hat{X} = s \cdot X_q$$
 (1)

where s is the scale factor, n is the bit-width,  $\lceil \cdot \rceil$  denotes round-to-nearest, and  $\operatorname{clip}()$  clamps values to  $[-2^{n-1}-1,\ 2^{n-1}-1]$ . Reducing bit-width provides near-linear storage compression, alleviating the growing memory demands of model parameters and KV-cache. At runtime, quantization also reduces memory traffic, thereby improving bandwidth efficiency. When activations are quantized in addition to weights, the core linear operation can be executed as an integer GEMM, further accelerating inference latency and throughput. For a transformer linear projection y = Wx,

$$y \approx s_W(W_q X_q) s_X \tag{2}$$

As widely recognized, the accuracy of LLM quantization for both weights and activations critically depends on how activation outliers are handled. Numerous approaches have been proposed to mitigate the quantization error introduced by such outliers. Early methods, such as SmoothQuant (Xiao et al. (2024)) and OmniQuant (Shao et al. (2024)), employ distribution-flattening techniques that balance activations and weights through pre-quantization scaling. Another line of work decomposes

matrix multiplication by assigning a separate high-precision path for outliers, as in LLM.int8() (Dettmers et al. (2022)) and QUIK (Ashkboos et al. (2023)). However, these primitive techniques exhibit significant accuracy degradation when applied to sub-8-bit LLMs.

More recently, online transformation techniques that rotate tensors to flatten the distribution have demonstrated effectiveness for 4-bit quantization with minimal latency overhead. Quarot (Ashkboos et al. (2024)) applies random Hadamard transformations, while SpinQuant (Liu et al. (2025)) learns rotation matrices to suppress outliers and reduce the performance variance. While effective in the W4A4 setting, existing methods still incur notable accuracy loss and either suffer from high variance due to random Hadamard matrices or require costly training to optimize transformation matrices, limiting deployment practicality. Meanwhile, LoRA-based LLM quantization has emerged as a powerful approach for mitigating quantization errors. In a similar vein, error reconstruction with low-rank factors via matrix decomposition has proven particularly effective, while avoiding heavy optimization or additional online layers. Further details are discussed in section 2.2.

#### 2.2 LOW-RANK ERROR RECONSTRUCTION

Recent advances in LoRA for parameter-efficient fine-tuning of foundation LLMs suggest its potential for applicability to LLM quantization as well. LoRA has been adapted to compensate for quantization errors by leveraging the auxiliary path of low-rank matrices, a strategy we refer to as low-rank error reconstruction. This approach introduces a LoRA-style compensator that restores the principal quantization error without requiring heavy retraining or complex deployment modifications. Concretely, it corrects the quantization error of a full-precision weight matrix  $W^{d\times d}$ , where d denotes the hidden dimension of a layer, by augmenting its low-bit proxy with low-rank factors.

$$\hat{\mathbf{W}} = \mathbf{Q}(\mathbf{W}) + \mathbf{L}_1 \mathbf{L}_2, \quad \mathbf{L}_1 \mathbf{L}_2 \approx \mathbf{W} - \mathbf{Q}(\mathbf{W})$$
(3)

Where  $Q(\boldsymbol{W})$  is the quantized weight,  $\boldsymbol{L}_1 \in \mathbb{R}^{d \times r}$ ,  $\boldsymbol{L}_2 \in \mathbb{R}^{r \times d}$ , and  $r \ll d$  is the rank. The parameter and compute overhead of the low-rank factors is only 2rd; with commonly used ranks  $r \in \{32, 64\}$  (e.g., d = 4096), this amounts to about 1.6-3.1%. Moreover, the LoRA-path is simple and lightweight, which makes it highly applicable across diverse model architectures and hardware platforms. Consequently, it has emerged as a prominent and widely adopted quantization method.

**Gradient-based Methods.** Gradient-based methods compensate the quantization errors by leveraging loss gradients on a small calibration set to learn low-rank factors, typically instantiating a compensator  $L_1L_2$  that reduces mismatch between the full-precision model and its quantized counterpart. QLoRA (Dettmers et al. (2023)) advanced this line of work by fine-tuning LoRA adapters on top of 4-bit weights, recovering task performance with modest resource overhead. At sub-4-bit precision, methods such as LQ-LoRA (Guo et al. (2024)), LoftQ (Li et al. (2023)) and QA-LoRA (Xu et al. (2023)) consistently mitigate degradation and sustain competitive accuracy, highlighting the viability of ultra-low-bit quantization.

**SVD-based Methods.** Another approach is to reconstruct quantization errors by low-rank matrices obtained via singular-value decomposition (SVD). This approach forms the compensator  $L_1L_2$  from the quantization error E by exploiting its rapidly decaying singular spectrum, allowing the dominant error to be captured with compact low-rank factors. Recent methods such as ZeroQuant-v2 (Yao et al. (2023)), CALDERA (Saha et al. (2024)), and  $L^2$ QER (Zhang et al. (2024)) adopt this SVD-based optimization to extract low-rank factors in a training-free and lightweight calibration process. In particular,  $L^2$ QER formalizes weight-activation quantization using SVD-based error reconstruction, preserving integer matrix multiplications across both the main and low-rank branches.

#### 2.3 Deployment of Low-Rank Error Reconstruction

In low-rank error reconstruction, a linear layer introduces an additional computation path for the low-rank branch, which cannot be parallelized cleanly with the main branch. Consider the case where both weights and activations are quantized. To enable the main branch to leverage low-precision matrix multiplication (e.g., INT4 or MXFP4 GEMM kernels), the LoRA-augmented linear layer must be evaluated in its decomposed form, such as  $X_q(W_q + L_1L_2) = X_qW_q + X_qL_1L_2$ . This requires two sequential matrix multiplications with low-rank matrices. Although each operation is relatively inexpensive, the overhead becomes non-negligible when the main branch benefits from

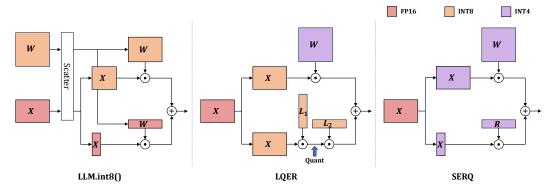


Figure 1: Computation flow of a linear layer under different matrix decomposition methods.  ${\tt LLM.int8}$  () employs a mixed-precision scheme of INT8 and FP16 by assigning separate computation paths for outliers and non-outliers.  $L^2QER$  applies SVD-based error reconstruction with a mixed-precision path of INT4 and INT8. In contrast, the proposed SERQ leverages a saliency-guided low-rank matrix and provides a unified computation path with INT4 or MXFP4 precision.

highly optimized low-precision GEMM kernels. To address this,  $L^2$ QER (Zhang et al. (2024)) quantizes the low-rank matrices as well, allowing the entire computation to proceed in a fully quantized path, as expressed in  $\boldsymbol{X}_q \boldsymbol{W}_q + \mathrm{Q}(\boldsymbol{X}_q \boldsymbol{L}_{1,q}) \boldsymbol{L}_{2,q}$ . However, this formulation introduces an additional on-the-fly quantization step for intermediate results between the sequential multiplications, resulting in inefficiency for low-precision deployment.

Figure 1 shows the overall datapath of prior approaches compared to ours, all based on decomposed matrix multiplication. Prior to error reconstruction methods, LLM.int8() (Dettmers et al. (2022)) handled outliers via a separate high-precision path, but this required on-the-fly scattering and FP16 computation, introducing substantial latency. An error reconstruction method  $L^2QER$  (Zhang et al. (2024)) later achieved a fully quantized datapath with mixed precision, yet still relied on on-the-fly quantization and two additional narrow matrix multiplications, resulting in latency and degraded accuracy in the W4A4 setting. In contrast, we propose SERQ, a novel error-reconstruction method that employs a single low-rank matrix, eliminating on-the-fly quantization and enabling a fully 4-bit end-to-end computation path. Further details are provided in section 3.

# 3 METHOD

We present SERQ, a saliency-aware error reconstruction method that jointly accounts for weight and activation saliency within a single low-rank matrix. We first show that saliency plays a central role in compensating weight-side quantization error, motivating a saliency-guided low-rank design. We then detail how SERQ incorporates activation saliency into the weight matrix, enabling error reconstruction based on integrated weight saliency while minimally impacting inference latency.

## 3.1 SALIENCY-AWARE LOW-RANK ADAPTATION

Prior error-reconstruction methods typically approximate the full-weight quantization error  $E = W - \mathrm{Q}(W)$  using a truncated SVD of the weight matrix,  $E \approx U_r \Sigma_r V_r^{\mathrm{T}}$ , where  $U_r \in \mathbb{R}^{d \times r}$  and  $V_r^{\mathrm{T}} \in \mathbb{R}^{r \times d}$  contain the top-r singular vectors and  $\Sigma_r \in \mathbb{R}^{r \times r}$  is the diagonal matrix of the corresponding singular values. While effective, this global decomposition overlooks where the error actually concentrates: the fixed rank budget is distributed across all rows and columns, diluting capacity on the most problematic regions (e.g., salient weights). In addition, quantizing the low-rank factors themselves introduces further loss, undermining the efficiency of error reconstruction.

Building on the theoretical insight that quantization errors vary across weight rows in linear operations, AWQ (Lin et al. (2024)) demonstrates that protecting only a small fraction ( $\sim$ 1%) of salient channels (i.e., their corresponding weight rows) using the activation distribution can substantially reduce the error. Motivated by this, we evaluate error reconstruction using SVD with a fixed rank size but restricted to salient rows selected by activation scales (see Appendix A.1). We find that selecting only a small number of salient rows for SVD improves perplexity compared to using the

Figure 2: (a) Overall SERQ implementation. During calibration, saliency rows are determined via activation scaling, followed by weight row permutation. During inference, error reconstruction is performed through a residual path computed only on the salient components, alongside the main path. (b) Computation flow of a decoder layer. The merged row- and column-wise weight permutation enables offline preprocessing of both current weight rows and subsequent activation channels.

entire matrix. In other words, rather than extracting ranks from the full weight matrix via SVD, we directly identify rank candidates based on row-wise saliency, allowing to reconstruct quantization error only on the most influential data. A narrowed matrix containing salient rows, where the row size equals the rank size, is sufficient for error reconstruction, enabling a single matrix decomposition for the low-rank branch. In section 3.2, we describe how to integrate activation statistics into the weight matrix to determine salient rows that disproportionately affect linear operations, and how to jointly reconstruct the error to improve accuracy using a full low-precision computing path.

# 3.2 IMPLEMENTATION

**Static Activation Flattening.** Activation quantization is notoriously fragile due to channel-wise outliers. Existing methods often rely on online outlier-handling techniques (e.g., rotation transforms or auxiliary layers), which are effective but introduce additional latency. Instead, since our method provides a residual path for error reconstruction at the linear layer, we avoid the online flattening process and revisit SmoothQuant (Xiao et al. (2024)), which flattens activation distributions using static per-channel scaling. Specifically, activations are scaled by a factor s, and the corresponding scale is folded into the weights. The operation in a linear layer can therefore be expressed as:

$$Y = XW = (X \cdot \operatorname{diag}(s^{-1}))(\operatorname{diag}(s) \cdot W) = \widetilde{X}\widetilde{W}$$
 (4)

The scaling factors are obtained during calibration and merged into adjacent layers offline, incurring no runtime overhead. While this strategy alleviates outliers, it shifts the quantization burden onto the weights, increasing the difficulty of weight quantization. However, unlike standalone prior approaches, the combined use of low-rank reconstruction enables effective compensation for the induced weight errors. Concretely, SERQ identifies salient weights after the flattening step and restores the residuals using a single low-rank matrix.

Saliency-Aware Error Reconstruction. The per-channel static flattening process pushes the scale of activation outliers into the corresponding weight rows. Assuming that the original weights follow a normal distribution, the salient rows in the folded weights can be identified directly by their scales. These rows subsequently accumulate significant errors when repeatedly multiplied with the activation matrix. To mitigate this, we introduce a low-rank compensator matrix  $R \in \mathbb{R}^{r \times d}$  that corrects the quantization errors in the r salient weight rows, denoted  $\widetilde{W}_s$ . Considering the weight

rows permuted in descending order of saliency (P), the folded matrix  $\widehat{\boldsymbol{W}}$  and the saliency-aware low-rank matrix  $\boldsymbol{R}$  can be defined as:

$$\widehat{\boldsymbol{W}} = P \cdot \operatorname{diag}(\boldsymbol{s}) \cdot \boldsymbol{W} = P \cdot \widetilde{\boldsymbol{W}} = [\widetilde{\boldsymbol{W}}_s; \widetilde{\boldsymbol{W}}_r], \quad \boldsymbol{R} = \widetilde{\boldsymbol{W}}_s - \operatorname{Q}(\widetilde{\boldsymbol{W}}_s)$$
 (5)

Here,  $W_r$  denotes the remaining weight rows. As shown, the residual errors of the salient rows are captured by the low-rank matrix to be used for reconstruction during matrix multiplication via an additional path. Then the overall linear operation can be described as:

$$Y = (X \cdot \operatorname{diag}(s) \cdot P^{-1})(P \cdot \operatorname{diag}(s) \cdot W) = \widehat{X}\widehat{W}$$
(6)

$$Q(\widehat{\boldsymbol{X}}) \cdot Q(\widehat{\boldsymbol{W}}) = Q([\widetilde{\boldsymbol{X}}_s \ \widetilde{\boldsymbol{X}}_r]) \cdot Q([\widetilde{\boldsymbol{W}}_s; \widetilde{\boldsymbol{W}}_r]) + Q(\widetilde{\boldsymbol{X}}_s) \cdot \boldsymbol{R}$$

$$\approx \widehat{\boldsymbol{X}}_q \cdot \widehat{\boldsymbol{W}}_q + \widetilde{\boldsymbol{X}}_{s,q} \cdot Q(\boldsymbol{R})$$
(7)

Importantly, we quantize the low-rank matrix as well, enabling pure low-precision computation across the entire path. Unlike SVD, directly extracting salient rows requires only a single matrix multiplication in the residual path, thereby eliminating intermediate quantization. Furthermore, only the activation channels corresponding to the salient rows  $(\widetilde{\boldsymbol{X}}_s \in \mathbb{R}^{s,len\times r})$  participate in the residual operation, enabling a computation-efficient low-rank multiplication of  $\mathbb{R}^{s\times r}\times\mathbb{R}^{r\times d}$ . The overall SERQ process, including both calibration and inference, is illustrated in Figure 1(a).

Offline Weight Permutation. We have aforementioned that the weights and activations must be properly reordered based on their saliency (e.g.  $\widehat{X} = [\widetilde{X}_s \ \widetilde{X}_r], \ \widehat{W} = [\widetilde{W}_s; \widetilde{W}_r]$ ). To address this, we propose a mergeable weight permutation scheme that eliminates latency overhead during inference. Both the rows and columns of the weight matrix are permuted offline according to the saliency order, enabling matrix multiplication to be executed directly on the appropriately reordered weights and activations. Figure 2(b) illustrates the computation flow of a single decoder layer, highlighting the offline permutation step. Based on saliency levels determined during calibration, row-wise permutations are prearranged for all weight parameters. The corresponding activations must follow the same channel order, which can be achieved by applying column-wise permutations to the preceding layer's weight matrix. For example, the permutation order  $P_4$  of the down-projection layer can be propagated to the weight columns of the preceding up- and gate-projection layers, ensuring that the activation outputs are produced in the desired order. Consequently, the resulting activations naturally align with  $P_4$ , allowing the down-projection to operate without additional reordering. In this way, all linear layers avoid on-the-fly reordering, and inference proceeds without latency overhead.

# 4 EXPERIMENTS

## 4.1 SETTINGS

Models and Tasks. We conduct experiments on LLaMA-2 (7B and 13B; Touvron et al. (2023)), LLaMA-3 (3.1 8B and 3.2 1B/3B; Grattafiori et al. (2024)), and Qwen-2.5 3B (Qwen et al. (2025)) models. Our evaluation covers eight zero-shot commonsense reasoning tasks—PIQA (Bisk et al. (2019)), SIQA (Sap et al. (2019)), ARC-Easy/Challenge (Clark et al. (2018)), HellaSwag (Zellers et al. (2019)), Winogrande (Sakaguchi et al. (2021)), BoolQ (Clark et al. (2019)), and OpenBookQA (Mihaylov et al. (2018))—and also provides perplexity scores on the WikiText2 test set (Merity et al. (2016)) as well as the MMLU benchmark (Hendrycks et al. (2021)). We further report results on generation tasks using the GSM8K (Cobbe et al. (2021)) and LongBench (Bai et al. (2024)) datasets.

Implementation Details. The calibration set is constructed from 128 random samples of the WikiText-2 dataset (Merity et al. (2016)) to identify salient rows. SERQ uses a rank size of 128, which is equivalent in effective bit width to using two low-rank matrices of rank 64. For quantization, the group size is set to 128, and weights are quantized using either GPTQ (Frantar et al. (2023)) or round-to-nearest (RTN) symmetric integer quantization, while activations are quantized with RTN asymmetric integer quantization. For the W4A4 configuration, we also provide the standard 4-bit Microscaling (MX) format, MXFP4, alongside the integer format, to demonstrate its implementation and measured speed on an NVIDIA RTX PRO 6000 GPU with Blackwell architecture support for MX format kernels (Rouhani et al. (2023)) (See Appendix A.2).

Table 1: Comparison with matrix decomposition methods. We compare perplexity scores, average zero-shot common sense reasoning accuracy, and average MMLU accuracy. Results under different precision settings are obtained by modifying their publicly released codebase (See Appendix A.5).

	_		-	-		-					
			l	LLaMA-2	7B	I	LaMA-2	13B	]	LLaMA-3	8B
#Bits	Method	<b>#Eff.</b> (w)	PPL(↓)	0-shot(↑)	MMLU(↑)	PPL(↓)	0-shot(↑)	MMLU(↑)	PPL(↓)	0-shot(↑)	MMLU(†)
FP16	baseline	16	5.47	64.09	41.83	4.88	66.53	52.04	6.13	67.16	62.13
	LLM.int4()	_†	7.28	59.64	29.01	5.35	65.01	47.12	10.82	60.86	46.72
W4A8	$L^2$ QER	4.35	5.83	63.35	39.4	5.08	66.39	50.98	7.16	65.68	57.81
W+A0	SERQ (RTN)	4.24	5.64	63.41	40.21	5	65.86	50.9	6.71	66.4	58.49
	SERQ (GPTQ)	4.24	5.59	63.04	40.29	4.98	65.66	49.46	6.52	66.23	60.25
	LLM.int4()	_†	6.32e+2	34.85	24.41	2.21e+3	34.74	23.04	4.87e+2	36.12	23.61
	$L^2$ QER	4.24	7.37	57.67	29.63	6.27	61.15	40.7	11.44	55.44	38.33
W4A4	L <sup>2</sup> QER-MXFP4	4.37	6.3	60.95	35.22	5.46	64.12	47.27	7.83	63.33	53.82
	SERQ (RTN)	4.24	6.03	61.77	38.03	5.24	65.14	49.22	8.07	62.49	51.84
	SERQ (GPTQ)	4.24	5.97	61.87	37.03	5.2	64.82	47.17	7.75	62.41	53.8
	SERQ-MXFP4	4.37	6.22	61.26	35.25	5.39	64.35	47.19	7.63	62.71	53.48
			L	LaMA-3.2	2 1B	L	LaMA-3.2	3B		Qwen-2.5-	3B
#Bits	Method	#Eff. (w)	PPL(↓)	0-shot(↑)	MMLU(↑)	PPL(↓)	0-shot(↑)	MMLU(↑)	PPL(↓)	0-shot(↑)	MMLU(†)
FP16	baseline	16	9.75	54.82	36.76	7.81	62.66	54.06	8.03	63.82	65.12
	LLM.int4()	_†	32.39	44.56	24.01	12.47	56.06	41.48	9.98e+5	35.59	25.52
W4A8	$L^2$ QER	4.35	12.06	50.97	31.23	8.78	60.36	51.24	12.25	54.71	26
W+A0	SERQ (RTN)	4.24	11.14	52.76	32.77	8.38	61.27	51.45	8.67	63.78	62.79
	SERQ (GPTQ)	4.24	10.45	53.93	27.1	8.18	61.54	50.24	8.36	63.7	63.25
	LLM.int4()	_†	1.7e+3	35.19	24.11	4.37e+2	35.77	23.5	8.51e+5	36.67	22.97
	$L^2$ QER	4.24	30.83	42.17	25.67	14.11	50.92	32.25	221.5	38.37	23.81
W4A4	L <sup>2</sup> QER-MXFP4		13.78	50.22	27.1	9.59	58.41	46.85	10.46	60.56	55.87
· · ¬/ 1·¬	SERQ (RTN)	4.24	13.57	50.15	29.11	9.43	58.4	46.36	9.66	61.24	59.52
	SERQ (GPTQ)	4.24	12.52	50.44	26.34	9.15	58.5	45.7	9.35	60.78	59.55
	SERQ-MXFP4	4.37	13.71	49.94	27.23	8.74	57.5	41.33	9.79	60.43	56.13

Compared Methods. We primarily compare our method against the state-of-the-art error reconstruction approach  $L^2QER$  under both W4A4 and W4A8 configurations. As a baseline for matrix decomposition with low-rank error reconstruction, we include LLM.int4(), the 4-bit variant of LLM.int8() (Dettmers et al. (2022)). We further compare against state-of-the-art W4A4 quantization methods that employ distribution-flattening techniques, including the rotation-based approaches Quarot Ashkboos et al. (2024) and SpinQuant Liu et al. (2025).

## 4.2 EVALUATION RESULTS

Comparison with Low-Rank Matrix Decomposition Methods. We evaluate the W4A8 and W4A4 precision configurations to demonstrate the effectiveness of SERQ in low-rank error reconstruction. Both  $L^2QER$  and SERQ are compared under the same group size, with low-rank matrix dimensions matched to ensure the same parameter counts. The difference in effective weight bit-width under W4A8 arises because  $L^2QER$  employs 8-bit precision for its low-rank matrices, whereas SERQ preserves 4-bit precision. As shown in Table 1, SERQ consistently outperforms LLM.int4() and  $L^2QER$  across most tasks, while achieving the lowest effective bit-width. Notably, SERQ is compatible with both RTN and GPTQ for weight quantization. Details of applying GPTQ robustly to SERQ are provided in Appendix A.3.

The accuracy gap is more pronounced under the W4A4 setting. While prior INT4 quantization methods suffer from severe degradation in most cases, SERQ maintains high accuracy across all tasks.  $L^2QER$  performs especially poorly on LLaMA-3 models, failing to preserve accuracy when both activations and low-rank matrices are quantized to 4-bit. We further provide MXFP4 implementations with the default group size of 32. Although MXFP4 proves adequate for both methods, the sequential reconstruction path in  $L^2QER$ , which requires two low-rank matrices, introduces significant latency overhead, which is examined in detail in the GPU performance analysis.

Comparison with W4A4 Distribution Flattening Methods. Rotation-based methods are recognized as state-of-the-art W4A4 quantization approaches, achieving the lowest effective bit-width with minimal additional parameters. We therefore compare two representative rotation methods, Quarot and SpinQuant, along with SmoothQuant's distribution-flattening baseline (Xiao et al. (2024)). The reported accuracies for rotation methods are obtained without key-value quantiza-

Table 2: Comparison with W4A4 distribution flattening methods. Latency overhead is measured as the additional computation time per linear layer relative to 4-bit GEMM (See Appendix A.5).

		#Eff.	Training-	Latency		LLaMA-2	7B	LLaMA-2 13B			
#Bits	Method	(w-bits)	free	overhead	PPL(↓)	0-shot(↑)	MMLU(↑)	PPL(↓)	0-shot(↑)	MMLU(↑)	
FP16	baseline	16	✓	81.4%	5.47	64.09	41.83	4.88	66.53	52.04	
	SmoothQ	~4	<b>V</b>	Х	1.51e+4	35.44	26.04	1.32e+4	34.53	23.85	
W4A4	SmoothQ(g128) OuaRot	4.13	/	<b>X</b> 19.8%	7.49 6.15	57.15 59.53	30.4 33.58	6.31 5.41	61.28 62.55	39.83 47.25	
	SpinQuant	~4	X	19.8%	6.0	61	34.8	5.2	64.8	47.8	
	SERQ-MXFP4	4.37		18.7%	6.22	61.26	35.25	5.39	64.35	47.19	
			LLaMA-38	BB		LLaMA-3.2	1B	LLaMA-3.2 3B			

			LLaMA-3	8B		LLaMA-3.2	1B	LLaMA-3.2 3B			
#Bits	Method	PPL(↓)	0-shot(↑)	MMLU(↑)	PPL(↓)	0-shot(↑)	MMLU(↑)	PPL(↓)	0-shot(↑)	MMLU(↑)	
FP16	baseline	6.13	67.16	62.13	9.75	54.82	36.76	7.81	62.66	54.06	
W4A4	SmoothQ SmoothQ(g128) QuaRot SpinQuant SERQ-MXFP4	4.71e+5 17.26 8.41 8.26 <b>7.63</b>	36.34 48.97 59.12 61.75 <b>62.71</b>	25.17 29.3 47.29 49.93 <b>53.48</b>	1.53e+5 69.22 <b>13.17</b> 13.47 13.71	35.59 40.04 <b>50.03</b> 48.95 49.94	24.37 24.43 26.64 26.38 <b>27.23</b>	3.73e+4 53.33 9.73 10.15 <b>8.74</b>	35.72 44.17 55.76 56.88 <b>57.5</b>	23.41 27.31 <b>44.75</b> 42.42 41.33	

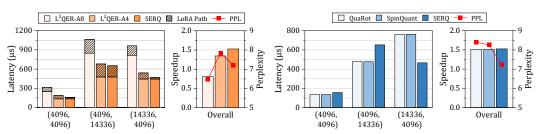


Figure 3: GPU performance comparison. Latency overhead introduced by the residual path is reported. SERQ is particularly effective for larger row-sized matrices (See Appendix A.4).

tion for fair comparison, using GPTQ without grouping for compatibility with low-precision GPU kernels. Consistently, although SERQ combined with GPTQ yields the best accuracy in our experiments, we primarily report results for its MXFP4 variant to ensure fair comparison under GPU kernels optimized for speed. As shown in Table 2, SpinQuant generally outperforms Quarot due to its learned rotation matrices, but this advantage diminishes on compact Llama-3.2 models, where both methods suffer from significant degradation. In contrast, SERQ achieves consistently higher accuracy, with the improvements most pronounced on Llama-3 models. Although SERQ incurs a slightly higher effective bit-width due to the inclusion of low-rank matrices and scaling factors in MX formatting, its per-layer latency overhead is lower than that of rotation-based methods.

**GPU Performance Comparison.** To examine inference performance with the proposed method, we measure execution time using low precision GEMM kernels implemented in the NVIDIA CUT-LASS library (cut (2025)). Since the fifth generation Blackwell Tensor Core architecture does not support INT4 GEMM kernels, while being substantially faster than the previous Ampere generation, we evaluate rotation-based methods in the FP4 format on the same latest core. Mixed precision computation for L<sup>2</sup>QER with W4A8 is measured with the MXFP mixed-precision kernel, while the Fast Hadamard transform for online rotation is executed using an FP32 kernel (Ashkboos et al. (2024)).

We benchmark the latency of linear operations across different dimension sizes in the LLaMA-3 8B decoder layer, and also report the overall speedup relative to an FP16 GEMM baseline. As shown in the left panel of Figure 3, SERQ outperforms both A4W4 and A4W8 settings of  $L^2QER$ , while maintaining competitive perplexity. Compared directly with  $L^2QER$  under A4W4, our single matrix error reconstruction path reduces latency overhead by up to  $4.5\times$  compared to the LoRA path, which requires two sequential low rank multiplications, thereby delivering the highest speedup. The right panel of Figure 3 compares performance against rotation-based methods. Since key-value quantization is excluded, the only online rotation occurs between the up or gate projection layer and the down projection layer. However, due to the unbalanced matrix dimensions, rotation introduces a significant overhead, about  $1.6\times$  greater than SERQ, which increases the latency of a single decoder layer. As a result, SERQ achieves the best perplexity score while delivering comparable speedups to rotation-based methods, with only about one percent additional latency overhead.

Table 3: Effect of rank size on perplexity.

#Rank	LLaMA-3 8B	LLaMA-3.2 1B	LLaMA-3.2 3B
0	9.8	15.09	10.21
16	8.28	14.32	9.71
32	8.24	14.09	9.59
64	8.18	13.97	9.57
128	8.07	13.57	9.43
256	7.98	13.25	9.32

Table 4: Effect of calibration data on perplexity.

Datasets	#Sample	LLaMA-38B	LLaMA-3.2 1B	LLaMA-3.2 3B
	512	8.07	13.57	9.43
Wiki	128	7.98	13.6	9.43
	32	8.08	13.51	9.45
	512	7.91	13.63	9.47
Pile	128	7.96	13.54	9.44
	32	8.18	13.44	9.52

Table 5: Generation task evaluation. SERQ is evaluated on GSM8K and LongBench datasets.

			GSM8K 5	i-shot(↑)		LongBench(1	`)
Model	#Bits	Method	flexible extract	strict match	qmsum	samsum	repobench-p
	FP16	baseline	48.07	47.69	23.43	46.26	66.56
LLaMA-3 8B	W4A8	LLM.int4() L <sup>2</sup> QER SERQ	7.88 36.24 <b>42.61</b>	4.62 35.33 <b>42.38</b>	13.29 7.28 <b>22.97</b>	34.2 25.81 <b>44.27</b>	44.93 47.03 <b>63.34</b>
	W4A4	L <sup>2</sup> QER SERQ	7.96 <b>23.65</b>	7.66 <b>23.12</b>	0.09 <b>19.08</b>	1.98 <b>40.85</b>	5.65 <b>54.78</b>
	FP16	baseline	26.08	25.93	23.94	42.98	64.42
LLaMA-3.2 3B	W4A8	LLM.int4() L <sup>2</sup> QER SERQ	6.76 18.95 <b>21.68</b>	6.44 18.42 <b>21.23</b>	15.83 21.53 <b>22.25</b>	36.42 40.63 <b>42.47</b>	35.24 60.18 <b>60.78</b>
	W4A4	L <sup>2</sup> QER SERQ	3.56 <b>16.22</b>	2.73 <b>15.77</b>	14.92 <b>19.31</b>	33.87 <b>42.14</b>	41.99 <b>53.67</b>

#### 4.3 ABLATION STUDIES

Sensitivity on the Rank Size. SERQ constructs a low rank matrix of size r. While we fix the rank size to align with prior methods, we also analyze its impact separately. We evaluate three models, LLaMA-3 8B and LLaMA-3.2 1B and 3B, by varying the rank size, which corresponds to the number of salient rows. As shown in Table 3, perplexity decreases monotonically with larger ranks, indicating improved accuracy. However, the improvement quickly saturates, and even the smallest setting of r=16 (equivalent to rank 8 in LoRA) yields competitive results.

**Sensitivity on Calibration Samples.** Table 4 reports perplexity scores obtained with varying calibration dataset sizes, which are used to determine saliency. We further evaluate our scheme on the Pile dataset Gao et al. (2020) to assess sensitivity to both sample size and dataset choice. The results indicate that SERQ is robust to the calibration dataset size and dataset characteristics, achieving similar perplexity scores across all settings.

**Evaluation on Generation Tasks.** While prior work on linear layer quantization primarily focuses on prefill-sensitive tasks (See section 4.1), we additionally evaluate generation tasks to demonstrate accuracy in generation. We conduct experiments on GSM8K (Cobbe et al. (2021)) and LongBench (Bai et al. (2024)) under W4A8 and W4A4 settings using LLaMA-3 8B and LLaMA-3.1 3B models. Compared with other matrix decomposition methods, SERQ achieves superior performance with only minor accuracy degradation in the W4A8 setting. Although the W4A4 setting leads to a notable accuracy drop overall, SERQ maintains reliable results where other methods fail.

## 5 CONCLUSION

In this work, we introduced SERQ, a saliency-aware error reconstruction method that enables end-to-end 4-bit quantization of both weights and activations using a single low rank matrix. By combining static activation flattening with mergeable weight permutation, SERQ identifies salient rows in the weight matrix without incurring additional latency overhead and reconstructs them through an auxiliary low rank branch. SERQ consistently outperforms prior matrix decomposition and distribution flattening methods, including state-of-the-art LoRA-based and rotation-based approaches, in terms of accuracy. Our implementation with the MXFP4 data format on NVIDIA Blackwell architecture further demonstrates that the auxiliary branch introduces minimal latency overhead, achieving significant speedups while preserving accuracy in low precision computation.

# REFERENCES

- Nvidia cutlass documentation. https://docs.nvidia.com/cutlass/index.html, 2025. Accessed: 2025-09-25.
- Saleh Ashkboos, Ilia Markov, Elias Frantar, Tingxuan Zhong, Xincheng Wang, Jie Ren, Torsten Hoefler, and Dan Alistarh. Quik: Towards end-to-end 4-bit inference on generative large language models, 2023. URL https://arxiv.org/abs/2310.09259.
- Saleh Ashkboos, Amirkeivan Mohtashami, Maximilian L. Croci, Bo Li, Pashmina Cameron, Martin Jaggi, Dan Alistarh, Torsten Hoefler, and James Hensman. Quarot: Outlier-free 4-bit inference in rotated llms, 2024. URL https://arxiv.org/abs/2404.00456.
- Yushi Bai, Xin Lv, Jiajie Zhang, Hongchang Lyu, Jiankai Tang, Zhidian Huang, Zhengxiao Du, Xiao Liu, Aohan Zeng, Lei Hou, Yuxiao Dong, Jie Tang, and Juanzi Li. Longbench: A bilingual, multitask benchmark for long context understanding, 2024. URL https://arxiv.org/abs/2308.14508.
- Yonatan Bisk, Rowan Zellers, Ronan Le Bras, Jianfeng Gao, and Yejin Choi. Piqa: Reasoning about physical commonsense in natural language, 2019. URL https://arxiv.org/abs/1911.11641.
- Christopher Clark, Kenton Lee, Ming-Wei Chang, Tom Kwiatkowski, Michael Collins, and Kristina Toutanova. Boolq: Exploring the surprising difficulty of natural yes/no questions, 2019. URL https://arxiv.org/abs/1905.10044.
- Peter Clark, Isaac Cowhey, Oren Etzioni, Tushar Khot, Ashish Sabharwal, Carissa Schoenick, and Oyvind Tafjord. Think you have solved question answering? try arc, the ai2 reasoning challenge, 2018. URL https://arxiv.org/abs/1803.05457.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. Training verifiers to solve math word problems, 2021. URL https://arxiv.org/abs/2110.14168.
- Tim Dettmers, Mike Lewis, Younes Belkada, and Luke Zettlemoyer. Llm.int8(): 8-bit matrix multiplication for transformers at scale, 2022. URL https://arxiv.org/abs/2208.07339.
- Tim Dettmers, Artidoro Pagnoni, Ari Holtzman, and Luke Zettlemoyer. Qlora: Efficient finetuning of quantized llms, 2023. URL https://arxiv.org/abs/2305.14314.
- Elias Frantar, Saleh Ashkboos, Torsten Hoefler, and Dan Alistarh. Gptq: Accurate post-training quantization for generative pre-trained transformers, 2023. URL https://arxiv.org/abs/2210.17323.
- Leo Gao, Stella Biderman, Sid Black, Laurence Golding, Travis Hoppe, Charles Foster, Jason Phang, Horace He, Anish Thite, Noa Nabeshima, Shawn Presser, and Connor Leahy. The pile: An 800gb dataset of diverse text for language modeling, 2020. URL https://arxiv.org/abs/2101.00027.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, Amy Yang, Angela Fan, et al. The llama 3 herd of models, 2024. URL https://arxiv.org/abs/2407.21783.
- Han Guo, Philip Greengard, Eric P. Xing, and Yoon Kim. Lq-lora: Low-rank plus quantized matrix decomposition for efficient language model finetuning, 2024. URL https://arxiv.org/abs/2311.12023.
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. Measuring massive multitask language understanding, 2021. URL https://arxiv.org/abs/2009.03300.
- Yixiao Li, Yifan Yu, Chen Liang, Pengcheng He, Nikos Karampatziakis, Weizhu Chen, and Tuo Zhao. Loftq: Lora-fine-tuning-aware quantization for large language models, 2023. URL https://arxiv.org/abs/2310.08659.

- Ji Lin, Jiaming Tang, Haotian Tang, Shang Yang, Wei-Ming Chen, Wei-Chen Wang, Guangxuan Xiao, Xingyu Dang, Chuang Gan, and Song Han. Awq: Activation-aware weight quantization for on-device llm compression and acceleration. In P. Gibbons, G. Pekhimenko, and C. De Sa (eds.), *Proceedings of Machine Learning and Systems*, volume 6, pp. 87–100, 2024. URL https://proceedings.mlsys.org/paper\_files/paper/2024/file/42a452cbafa9dd64e9ba4aa95cclef21-Paper-Conference.pdf.
- Zechun Liu, Changsheng Zhao, Igor Fedorov, Bilge Soran, Dhruv Choudhary, Raghuraman Krishnamoorthi, Vikas Chandra, Yuandong Tian, and Tijmen Blankevoort. Spinquant: Llm quantization with learned rotations, 2025. URL https://arxiv.org/abs/2405.16406.
- Stephen Merity, Caiming Xiong, James Bradbury, and Richard Socher. Pointer sentinel mixture models, 2016. URL https://arxiv.org/abs/1609.07843.
- Todor Mihaylov, Peter Clark, Tushar Khot, and Ashish Sabharwal. Can a suit of armor conduct electricity? a new dataset for open book question answering, 2018. URL https://arxiv.org/abs/1809.02789.
- Markus Nagel, Marios Fournarakis, Rana Ali Amjad, Yelysei Bondarenko, Mart van Baalen, and Tijmen Blankevoort. A white paper on neural network quantization. *CoRR*, abs/2106.08295, 2021. URL https://arxiv.org/abs/2106.08295.
- Qwen, :, An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, Huan Lin, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiaxi Yang, Jingren Zhou, Junyang Lin, Kai Dang, Keming Lu, Keqin Bao, Kexin Yang, Le Yu, Mei Li, Mingfeng Xue, Pei Zhang, Qin Zhu, Rui Men, Runji Lin, Tianhao Li, Tianyi Tang, Tingyu Xia, Xingzhang Ren, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yu Wan, Yuqiong Liu, Zeyu Cui, Zhenru Zhang, and Zihan Qiu. Qwen2.5 technical report, 2025. URL https://arxiv.org/abs/2412.15115.
- Bita Darvish Rouhani, Ritchie Zhao, Ankit More, Mathew Hall, Alireza Khodamoradi, Summer Deng, Dhruv Choudhary, Marius Cornea, Eric Dellinger, Kristof Denolf, Stosic Dusan, Venmugil Elango, Maximilian Golub, Alexander Heinecke, Phil James-Roxby, Dharmesh Jani, Gaurav Kolhe, Martin Langhammer, Ada Li, Levi Melnick, Maral Mesmakhosroshahi, Andres Rodriguez, Michael Schulte, Rasoul Shafipour, Lei Shao, Michael Siu, Pradeep Dubey, Paulius Micikevicius, Maxim Naumov, Colin Verrilli, Ralph Wittig, Doug Burger, and Eric Chung. Microscaling data formats for deep learning, 2023. URL https://arxiv.org/abs/2310.10537.
- Rajarshi Saha, Naomi Sagan, Varun Srivastava, Andrea J. Goldsmith, and Mert Pilanci. Compressing large language models using low rank and low precision decomposition, 2024. URL https://arxiv.org/abs/2405.18886.
- Keisuke Sakaguchi, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. Winogrande: an adversarial winograd schema challenge at scale. *Commun. ACM*, 64(9):99–106, August 2021. ISSN 0001-0782. doi: 10.1145/3474381. URL https://doi.org/10.1145/3474381.
- Maarten Sap, Hannah Rashkin, Derek Chen, Ronan LeBras, and Yejin Choi. Socialiqa: Commonsense reasoning about social interactions, 2019. URL https://arxiv.org/abs/1904.09728.
- Wenqi Shao, Mengzhao Chen, Zhaoyang Zhang, Peng Xu, Lirui Zhao, Zhiqian Li, Kaipeng Zhang, Peng Gao, Yu Qiao, and Ping Luo. Omniquant: Omnidirectionally calibrated quantization for large language models, 2024. URL https://arxiv.org/abs/2308.13137.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, et al. Llama 2: Open foundation and fine-tuned chat models, 2023. URL https://arxiv.org/abs/2307.09288.
- Guangxuan Xiao, Ji Lin, Mickael Seznec, Hao Wu, Julien Demouth, and Song Han. Smoothquant: Accurate and efficient post-training quantization for large language models, 2024. URL https://arxiv.org/abs/2211.10438.

Yuhui Xu, Lingxi Xie, Xiaotao Gu, Xin Chen, Heng Chang, Hengheng Zhang, Zhengsu Chen, Xi-aopeng Zhang, and Qi Tian. Qa-lora: Quantization-aware low-rank adaptation of large language models, 2023. URL https://arxiv.org/abs/2309.14717. Zhewei Yao, Xiaoxia Wu, Cheng Li, Stephen Youn, and Yuxiong He. Zeroquant-v2: Exploring post-training quantization in llms from comprehensive study to low rank compensation, 2023. URL https://arxiv.org/abs/2303.08302. Rowan Zellers, Ari Holtzman, Yonatan Bisk, Ali Farhadi, and Yejin Choi. Hellaswag: Can a ma-chine really finish your sentence?, 2019. URL https://arxiv.org/abs/1905.07830. Cheng Zhang, Jianyi Cheng, George A. Constantinides, and Yiren Zhao. Lqer: Low-rank quantiza-tion error reconstruction for llms, 2024. URL https://arxiv.org/abs/2402.02446. Weibo Zhao, Yubin Shi, Xinyu Lyu, Wanchen Sui, Shen Li, and Yong Li. Aser: Activa-tion smoothing and error reconstruction for large language model quantization, 2024. URL https://arxiv.org/abs/2411.07762. 

## A APPENDIX

#### A.1 EXPERIMENTS ON LOW-RANK SALIENCY

In this analysis, we vary the number of weight rows included in SVD, selected by saliency in descending order of activation channel scales, to examine the trade-off between loss from rank reduction and the coverage of error reconstruction. To isolate the quantization effects of the weight matrix, we apply W4A16 quantization with a group size of 128 and fix the SVD rank size to 64. Figure 4 shows that decomposing only salient weights is more effective for reducing quantization error, as the perplexity score decreases slightly on the leftmost side when using row size 64. This result indicates that restricting reconstruction to salient rows not only avoids accuracy degradation but also yields consistent improvement across all tested LLaMA models. We observe about 1 to 4 percent improvements compared to using the same limited rank to cover the full weight matrix. Expanding beyond the salient subset offers little or even negative benefit, since the additional low rank factorization loss outweighs the marginal coverage. These findings motivate a single-layer reconstruction, rather than two low-rank factors, and support a saliency-aware error reconstruction strategy that keeps both the main and auxiliary paths in pure 4-bit, achieving efficiency without sacrificing accuracy.

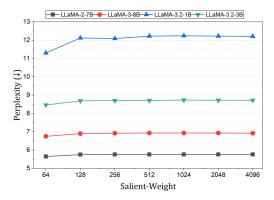


Figure 4: The trade-off between loss from rank reduction and the coverage of error reconstruction. The figure shows that higher accuracy is achieved by reconstructing errors for salient rows with smaller ranks, rather than covering a larger portion of the weight matrix.

## A.2 MICROSCALING (MX) FORMAT

The Microscaling (MX) format Rouhani et al. (2023) is a block-scaled, low-precision representation that associates each small block with a single scale and quantized elements, thereby retaining dynamic range while enabling efficient 4–8 bit computation. Concretely, data are partitioned into blocks; each block stores one scale, estimated from its maximum value, and the elements within the block are stored in low-bit precision—either in floating-point formats (MXFP8, MXFP6, MXFP4) or in integer format (MXINT8). Note that INT4 is not included among MX variants. All elements are defined relative to the block's shared scale, which is fixed at 8-bit precision.

In contrast, naive single-scale low-bit formats often degrade sharply on LLMs: a few outliers inflate the global scale, wasting quantization levels and clipping inlier values. By adopting a blocklevel granularity (typically set to 32), MX preserves usable dynamic range and maintains effective quantization. As a result, MX-based models achieve low-precision inference with minimal accuracy loss, reduced memory and bandwidth, and hardware-friendly execution, since computations reduce to standard low-bit arithmetic plus lightweight per-block scaling. Owing to its effectiveness for low-precision quantization in deep learning, NVIDIA's Blackwell architecture, equipped with fifth-generation Tensor Cores, provides CUTLASS kernels supporting MX data formats for LLM workloads.

# A.3 GPTQ IMPLEMENTATION

Recently GPTQ has shown strong performance in weight-only quantization. It performs a lightweight, Hessian-guided optimization: quantizing weights along principal curvature directions and allowing the remaining (unquantized) weights to compensate the induced error. However, this procedure is not independent of error-reconstruction pipelines. As GPTQ updates proceed, the compensation steps shift the weight distribution, so the naïve error-reconstruction modeling  $E = W - Q(W) = U\Sigma V^T$ —i.e., using a fixed pre-quantized weight and a post-quantized weight—no longer yields a clean comparison, making the simple SVD-based approximation less well-aligned and less effective. before quantization to apply GPTQ effectively in the weight matrix. Specifically, we substitute the low rank matrix R into the salient weight rows  $\widetilde{W}_s$  prior to the iterative GPTQ process, where  $R = Q(\widetilde{W}_s)$  and the salient rows are assigned the error  $\widetilde{W}_s - Q(\widetilde{W}_s)$ . This substitution diverts activation-driven outlier effects away from the main weights and stabilizes the weight distribution. Afterward, GPTQ on the main path yields higher error reconstruction efficiency and simplifies implementation by avoiding dequantization and requantization cycles when building the error reconstruction term. In this way, GPTQ targets a better conditioned main weight, while the low rank path captures the salient components.

## A.4 GPU PERFORMANCE ANALYSIS DETAILS

This section reports the absolute latency obtained with NVIDIA Blackwell CUTLASS kernels when operating on linear layers of different sizes used in LLM models. We also provide the inference overhead introduced by low rank factors, where SERQ includes an additional multiplication path using R, and  $L^2$ QER includes the factors  $L_1$  and  $L_2$ .

Table 6: GPU latency results in linear layers.

				Inf	erence ove	rhead (MX	KFP4)
Sequence Length	Weight Size	$\mathbf{FP16}(\mu s)$	$\mathbf{MXFP4}(\mu \mathbf{s})$	$R(\mu s)$	$\boldsymbol{L}_1(\mu \mathrm{s})$	$\boldsymbol{L}_2(\mu \mathrm{s})$	$\boldsymbol{L}_1\boldsymbol{L}_2(\mu s)$
	$2048 \times 2048$	71.402	27.457	10.501	16.606	12.009	28.615
	2048 × 8192	207.482	72.7	20.934	16.606	24.86	41.466
	8192 × 2048	212.87	88.738	10.501	43.235	12.009	55.244
2048	4096 × 4096	211.894	70.255	14.637	26.833	14.66	41.493
	4096 × 11008	519.322	181.359	27.0112	26.833	31.034	57.867
	$11008 \times 4096$	512.25	175.964	14.637	55.51	14.66	70.171
	4096 × 14336	643.27	235.397	38.342	26.833	40.638	67.472
	$\overline{14336 \times 4096}$	655.757	226.379	14.637	69.887	14.659	84.547
	8192 × 8192	755.494	260.054	20.934	43.235	24.86	68.095
	$8192 \times 28672$	2604.973	946.403	174.537	43.235	174.857	218.09
	$28672 \times 8192$	2444.912	946.787	20.934	132.005	24.86	156.865
	$2048 \times 2048$	122.298	37.664	14.4035	16.669	14.662	31.331
	2048 × 8192	403.859	142.159	56.211	16.669	55.991	72.66
	$8192 \times 2048$	394.138	132.903	14.4035	43.292	14.662	57.954
4096	$4096 \times 4096$	382.79	136.076	20.936	26.805	25.155	51.96
	4096 × 11008	1009.898	355.986	129.704	26.8045	129.876	156.681
	$11008 \times 4096$	986.624	345.739	20.938	55.947	25.155	81.1
	4096 × 14336	1262.093	477.976	175.101	26.8045	175.268	202.073
	$14336 \times 4096$	986.97	445.022	20.938	70.016	25.155	95.171
	8192 × 8192	1560.854	506.364	56.211	43.292	55.991	99.283
	8192 × 28672	4868.259	1911.05	355.274	43.292	355.969	399.261
	$28672 \times 8192$	4849.363	1853.27	56.211	132.019	55.991	188.01

## A.5 OVERALL LLM ACCURACY RESULTS

We report overall accuracy results on the evaluated benchmarks, which are not captured in Tables 1 and 2 since they only present average values. Tables 7 and 8 provide detailed accuracy results for low-rank matrix decomposition methods under the W4A8 and W4A4 settings, respectively. Table 9 reports detailed results for different distribution flattening methods.

Table 7: Accuracy results of low-rank matrix decomposition methods when tested with W4A8 settings.

					0-Sho	t Comm	on Sense	Reasoni	ng tasks					MMLU		
Model	Methods	PPL (↓)	BoolQ (†)	PIQA (†)	SIQA (†)	ARC-e (↑)	ARC-c	HellaS.	WinoG.	OBQA (†)	Avg. (†)	Human.	Other	SocialS.	STEM (†)	<b>Avg.</b> (↑)
	baseline	5.47	77.74	79.05	46.11	74.49	46.25	76	68.9	44.2	64.09	39.72	47.12	47.45	34.28	41.83
LLaMA-2-7B	LLM.int4()	7.28	69.3	76.39	43.76	68.48	43.43	71.19	64.96	39.6	59.64	26.23	27.9	33.54	29.84	29.01
ELMINIA 2 7B	$L^2$ QER	5.83	75.99	78.51	44.98	74.75	44.03	75.14	68.98	44.4	63.35	37.39	43.13	43.48	34.76	39.4
	SERQ (RTN) SERQ (GPTQ)	5.64 5.59	76.85 77	78.4 78.78	44.93 44.83	73.7 73.15	45.31 44.2	75.12 75.05	68.75 68.27	44.2 43	63.41 63.04	37.39 37.9	45.93 45.22	45.47 44.85	33.65 34.54	40.21 40.29
					0-Sho	t Comm	on Sense	Reasoni	ng tasks					MMLU		
Model	Methods	PPL	BoolQ	PIQA	SIQA	ARC-e	ARC-c	HellaS.	WinoG.	OBQA	Avg.	Human.	Other	SocialS.	STEM	Avg.
Wiodei	Wiethous	(\psi)	(†)	(†)	(†)	(†)	(†)	(†)	(↑)	(†)	(↑)	(↑)	(†)	(†)	(†)	(†)
	baseline	4.88	80.61	80.52	47.39	77.53	49.23	79.38	72.38	45.2	66.53	47.89	59.29	61.16	42.21	52.04
LLaMA-2-13B	LLM.int4()	5.35	79.91	79.38	46.52	75.34	49.15	77.46	69.69	42.6	65.01	44.31	53.07	54.76	38	47.12
	L <sup>2</sup> QER SERQ (RTN)	5.08	81.41 80.18	79.87 80.36	46.37 46.47	76.85 76.3	50.17 48.38	78.52 78.7	72.3 71.67	45.6 44.8	66.39 65.86	46.63 47.08	58.71 57.64	59.02 58.99	42.02 42.09	50.98 50.9
	SERQ (RTN) SERQ (GPTQ)	4.98	79.54	79.38	46.26	76.18	47.27	78.83	72.45	45.4	65.66	45.59	55.52	58.01	40.91	49.46
					0-Sho	t Comm	on Sense	Reasoni	ng tasks					MMLU		
Model	Methods	PPL	BoolQ		SIQA				WinoG.		Avg.	Human.			STEM	
	Less Park	(↓)	(↑)	(↑)	(↑)	(↑)	(↑)	79.2	(↑)	(↑)	(↑)	(↑)	(↑)	(↑)	(↑)	(†)
	baseline	6.13	81.07	80.74	47.08	77.69	52.99		73.48	45	67.16	54.81	70.55	73.25	53.89	62.13
LLaMA-3-8B	LLM.int4() L <sup>2</sup> QER	10.82 7.16	75.23 80.67	74.97 79.54	44.22 45.29	67.38 75.34	42.24 51.28	72.63 77.01	70.01 72.53	40.2 43.8	60.86 65.68	42.34 52.01	53.91 66.4	52.65 67.66	40.41 48.37	46.72 57.81
	SERQ (RTN)	6.71	81.41	80.2	45.29	77.44	50.77	77.61	73.64	44.4	66.4	51.9	67.94	69.16	48.59	58.49
	SERQ (GPTQ)	6.52	80.8	80.9	46.21	76.52	51.02	78.25	72.53	43.6	66.23	53.07	69.46	70.62	51.76	60.25
					0-Sho	t Comm	on Sense	Reasoni	ng tasks					MMLU		
Model	Methods	PPL		PIQA					WinoG.	_	Avg.	Human.		SocialS.		
		(↓)	(↑)	(↑)	(↑)	(↑)	(↑)	(↑)	(†)	(†)	(↑)	(↑)	(↑)	(↑)	(†)	<u>(↑)</u>
	baseline	9.75	63.67	74.48	42.84	60.44	36.18	63.73	59.83	37.4	54.82	34.92	41.1	39.78	32.29	36.76
LLaMA-3.2-1B	LLM.int4()	32.39	51.07	63.82	38.74	42.09	28.24	46.89	55.25	30.4	44.56	25.1	24.24	23.17	22.96	24.01
	L <sup>2</sup> QER SERQ (RTN)	12.06 11.14	60.67	70.73	41.45	53.45	32.68	59.04	56.91 56.35	32.8	50.97	28.59	35.92	33.99	27.85 28.7	31.23
	SERQ (RTR) SERQ (GPTQ)	10.45	63.24	72.47	42.02	58.92	36.95	61.7	58.96	37.2	53.93	27.72	29	27.72	23.69	27.1
			14 62.75 72.85 42.07 56.69 34.39 60.74 56.35 36.2 52.76 32.56 36.53 33.44 2													
Model	Methods	PPL	_	PIQA	_	ARC-e	ARC-c	HellaS.		-		Human.	Other	SocialS.		
		(\psi)	(†)	(†)	(†)	(†)	(†)	(†)	(↑)	(†)	(↑)	(↑)	(†)	(†)	(†)	(†)
	baseline	7.81	72.97	77.53	46.93	71.59	46.25	73.49	69.53	43	62.66	48.78	63.08	62.59	44.72	54.06
LLaMA-3.2-3B	LLM.int4()	12.47	67.55	73.99	42.94	62.58	37.46	66.57	61.8	35.6	56.06	37.26	46.6	47.71	36.63	41.48
	L <sup>2</sup> QER SERQ (RTN)	8.78 8.38	71.8 71.22	76.61 76.55	45.8 46.16	66.29 69.19	41.47 44.28	71.13 72.31	67.8 69.85	42 40.6	60.36	46.93 46.91	58.32 59.86	59.67 59.02	42.47 42.53	51.24 51.45
	SERQ (RTN) SERQ (GPTQ)	8.18	71.19	76.44	46.83	68.9	45.39	72.31	69.14	42	61.54	46.76	58.93	56.65	40.63	50.24
	-				0-Sho	t Comm	on Sense	Reasoni	ng tasks					MMLU		
Model	Methods	PPL	BoolQ	PIQA	SIQA	ARC-e	ARC-c	HellaS.	WinoG.	OBQA	Avg.	Human.	Other	SocialS.	STEM	Avg.
Model	Methods	(\dul)	(†)	( )	(†)	(↑)	(↑)	(↑)	(↑)	(↑)	(†)	(↑)	(†)	(↑)	(↑)	(↑)
	baseline	8.03	77.58	78.73	50.1	73.11	47.1	73.62	68.51	41.8	63.82	56.77	71.07	76.15	60.96	65.12
Qwen-2.5 3B	LLM.int4()	9.98e+5	37.86	51.52	34.19	24.58	26.19	26.51	51.3	32.6	35.59	27.14	24.01	23.79	26.26	25.52
-	L <sup>2</sup> QER	12.25	68.1	73.07	42.99	47.22	37.54	66.57	64.01	38.2	54.71	26.82	26.97	25.22	24.58	26
	SERQ (RTN) SERQ (GPTQ)	8.67 8.36	76.85 77.68	77.75 77.86	50.41 49.23	76.68 75.46	48.38 48.63	72.13 72.35	65.82 66.61	42.2 41.8	63.78 63.7	55.39 55.07	68.46 69.46	73.71 74.65	57.6 58.2	62.79 63.25

## A.6 LLM USAGE

 In this paper, we used a large language model to aid and polish the manuscript to improve the overall writing quality.

Table 8: Accuracy results of low-rank matrix decomposition methods when tested with W4A4 settings.

Methods						0-Sho	t Comm	on Sense	Reasoni	ng tasks					MMLU		
Liama   Liam	Model	Methods															Avg. (↑)
1,2   QER   1,3   3,6   3,7   3,6   3,7   3,6   3,8   4,90   6,95   3,		baseline	5.47	77.74	79.05	46.11	74.49	46.25	76	68.9	44.2	64.09	39.72	47.12	47.45	34.28	41.83
1	LLaMA-2-7B		6.32e+2	43.3	49.84	34.08	27.44	25.43	26.46	50.28	22	34.85	25.04	23.46	24.6	24.2	24.41
Seric (Seric (Seric (Seric (Seric (Seric ) 1																	29.63
SERQ (CPT)   5,97   74,56   78,13   34.2   72,9   43,09   73,88   66,61   41   61,68   74,67   41,31   40,91   31,43   35,24   37,8																	38.03
Methods		SERQ (GPTQ)	5.97	74.56	78.13	44.22	72.9	43.09	73.88	66.61		61.87	34.67	41.33	40.92	32.54	37.03
Mode    Methods   PPL   Mode    PPL		SERQ-MAFP4	0.22		11.2						41	01.20	32.10			31.43	33.23
Nethods			PPI.	Rool()	PIOA						OROA	Δνα	Human			STEM	Δνσ
Liama-2-138	Model	Methods															(†)
L2    C2  C2  C3  C3  C3  C4  C3  C5  C4  C3  C5		baseline	4.88	80.61	80.52	47.39	77.53	49.23	79.38	72.38	45.2	66.53	47.89	59.29	61.16	42.21	52.04
1	LLaMA-2-13B																23.04
SERQ (RTN)   5.24   78.35   77.36   76.39   46.42   78.11   70.64   44.6   64.22   43.83   53.3   53.45   83.42   71.85   77.86   77.86   78.79   78.70   78																	47.27
SERQ-MXFP4   5.39   77.86   78.79   45.19   75.59   48.12   76.18   70.09   43   64.35   43.59   52.33   54.6   40.28   47.14		SERQ (RTN)	5.24	78.35	80.25	46.37	76.39	46.42	78.11	70.64	44.6	65.14	45.89	55.29	56.45	41.14	49.22
Mode    Methods   PPL   Mode    Mode    Methods   PPL   Mode																	47.17 47.19
Methods										ng tasks					MMLU		
	Model	Methods		-	-												
LLaMA-3.8-B   LLMint=4 ()																	(†) (2.12
L <sup>2</sup> QER   11.44   61.71   72.63   42.02   62.37   39.59   67.29   60.54   37.4   55.44   35.47   42.68   42.7   34.06   38.3     L <sup>2</sup> QER   11.44   61.71   72.63   42.01   74.64   74.64   69.38   41.6   69.38   41.6   69.33   41.6   60.22   63.29   45.8   53.8   51.8     SERQ (RTN)   8.07   74.5   76.06   44.37   71.93   46.67   72.98   70.64   42.8   62.49   46.14   60.22   60.22   43.83   51.8     SERQ (MTNP   7.63   76.15   77.2   44.11   72.69   44.88   73.6   69.38   40.8   62.41   48.42   61.6   62.2   43.29   53.8     Model   Methods   PPL   BoolQ   PIQA   SIQA   ARC-e   ARC-e   HellaS.   WinoG   OBQA   Avg.   Human   Other   SocialS.   STEM   Avg.     LLaMA-3.2-1B   LLM.int4 (1)   1.7e+3   41.13   51.03   33.73   26.89   22.61   26.92   51.38   27.8   35.19   24.21   24.59   24.47   23.15   24.1     L <sup>2</sup> QER   30.83   51.19   60.01   36.18   40.95   26.62   43.41   52.57   26.4   42.17   25.02   26.68   25.41   23.88   25.6     SERQ (RTN)   13.78   61.44   69.24   62.23   53.91   30.03   55.04   56.67   32.2   49.94   27.55   28.23   27.49   25.5   27.2     Model   Methods   PPL   BoolQ   PIQA   SIQA   ARC-e   ARC-e   HellaS.   WinoG   OBQA   Avg.   Human   Other   SocialS.   STEM   Avg.     LLaMA-3.2-1B   LLM.int4 (1)   1.7e+3   41.13   51.03   33.73   26.89   22.61   26.92   51.38   27.8   35.19   24.21   24.59   24.47   23.15   24.1     L <sup>2</sup> QER   31.78   61.44   69.24   62.23   53.91   30.03   55.04   56.75   33.0   50.22   26.68   25.46   25.15   27.2     Model   Methods   PPL   BoolQ   PIQA   SIQA   ARC-e   ARC-e   HellaS.   WinoG   OBQA   Avg.   Human   Other   SocialS.   STEM   Avg.   Avg.	11 14 2 00																
SERQ (RTN)   SO, 7 4.5   76.06   44.37   71.93   46.67   72.98   70.08   42.8   62.49   46.14   60.02   60.32   43.83   51.8	LLaMA-3-8B																38.33
SERQ (GPTQ)   7.75   76.85   77.8   43.55   72.43   44.88   73.6   69.48   40.8   62.41   48.42   61.6   62.2   45.92   53.5																	53.82
Model   Methods   Pri																	51.84
Methods   PPL   BoolQ   PIQA   SIQA   ARC-e   ARC-c   HellaS   WinoG.   OBQA   Avg.   Human.   Other   Othe																	53.48
Methods   (↓) (↑) (↑) (↑) (↑) (↑) (↑) (↑) (↑) (↑) (↑						0-Sho	t Comm	on Sense	Reasoni	ng tasks					MMLU		
LLaMA-3.2-18   LLM_int4()   1.7e+3   41.13   51.03   33.73   26.89   22.61   26.92   51.38   27.8   35.19   24.21   24.59   24.47   23.15   24.1	Model	Methods															<b>Avg.</b> (↑)
L <sup>2</sup> QER   30.83   51.19   60.01   36.18   40.95   26.62   43.41   52.57   26.4   42.17   25.02   26.68   25.41   25.88   25.68   L <sup>2</sup> QER-MXFP4   13.78   61.44   69.26   42.32   53.91   30.03   55.04   56.75   33   50.22   26.67   28.68   26.86   25.15   27.		baseline	9.75	63.67	74.48	42.84	60.44	36.18	63.73	59.83	37.4	54.82	34.92	41.1	39.78	32.29	36.76
L2   QER-MXFP4   13.78   61.44   69.26   42.32   53.91   30.03   55.04   56.75   33.   50.22   26.67   28.68   26.86   25.15   27.	LLaMA-3.2-1B																24.11
SERQ (RTN)   13.57   57.68   69.26   40.28   55.93   32.94   56.63   55.25   33.2   50.15   28.78   30.51   29.64   27.69   29.1																	25.67
SERQ (GPTQ   12.52   59.85   70.13   40.58   53.41   32.68   57.84   54.62   34.4   50.44   26.93   27.84   26.13   24.2   26.3   27.49   25.5   27.2																	29.11
Model   Methods   PPL   Model   PPL   Mod		SERQ (GPTQ)	12.52	59.85	70.13	40.58	53.41	32.68	57.84	54.62	34.4	50.44	26.93	27.84	26.13		26.34
Methods   PPL   BoolQ   PIQA   SIQA   ARC-e   ARC-c   Hellas   WinoG.   OBQA   Avg.   Human.   Other   SocialS   STEM   Avg.		SERQ-MXFP4	13./1	60.03	68.//						32.2	49.94	27.55			25.5	21.23
Methods			DDI		DIOA						OBOA	<b>A</b>				CTEM	
$ \begin{array}{c ccccccccccccccccccccccccccccccccccc$	Model	Methods															(↑)
$ \begin{array}{c ccccccccccccccccccccccccccccccccccc$		baseline	7.81	72.97	77.53	46.93	71.59	46.25	73.49	69.53	43	62.66	48.78	63.08	62.59	44.72	54.06
$ \begin{array}{c ccccccccccccccccccccccccccccccccccc$	LLaMA-3.2-3B		4.37e+2	45.41	52.67	33.93	28.07	22.61	27.94	49.49	26	35.77	25.02	24.27	22.07	21.88	23.5
SERQ (RTN)   9.43   67.74   73.72   45.45   66.33   42.32   69.09   64.96   37.6   58.4   43.06   53.91   51.09   39.23   46.2     SERQ (GPTQ)   9.15   67.52   74.48   45.29   64.73   42.06   69.57   66.38   38   58.5   42.27   53.97   50.37   38.09   45.2     SERQ-MXFP4   8.74   68.56   73.83   45.24   61.83   40.02   671.6   64.33   39   57.5   38.68   47.89   45.82   34.44   41.3     Model   Methods   PPL   BoolQ   PIQA   SIQA   ARC-e   ARC-c   HellaS   WinoG   OBQA   Avg   Human   Other   SocialS   STEM   Avg     Mathods   Serger																	32.25
SERQ (GPTQ)   9.15   67.52   74.48   45.29   64.73   42.06   69.57   66.38   38   58.5   42.27   53.97   50.37   38.09   45.5																	46.85 46.36
Model   Methods   PPL   Bool   PIQA   SIQA   ARC-e   Melas		SERQ (GPTQ)	9.15	67.52	74.48	45.29	64.73	42.06	69.57	66.38	38	58.5	42.27	53.97	50.37	38.09	45.7
$ \begin{array}{c ccccccccccccccccccccccccccccccccccc$		SERQ-MXFP4	8.74	08.30	13.83						39	37.3	38.08			34.44	41.33
$ \begin{array}{ c c c c c c c c c c c c c c c c c c c$	M. 1.1	M.d. J	PPL	BoolO	PIOA						OBOA	Avg.	Human.			STEM	Avg.
$ \begin{array}{c ccccccccccccccccccccccccccccccccccc$	Model	ivietnods															(↑)
$ \begin{array}{c ccccccccccccccccccccccccccccccccccc$			8.03														65.12
$ \begin{array}{cccccccccccccccccccccccccccccccccccc$	Qwen-2.5-3B																22.97
SERQ (RTN) 9.66 72.6 75.63 49.03 72.14 44.11 70 64.64 41.8 61.24 52.52 65.88 69.97 53.5 59.5 SERQ (GPTQ) 9.35 74.74 75.73 45.55 72.14 45.31 69.92 63.22 39.6 60.78 52.54 65.92 70.43 53.12 59.5																	23.81
SERQ (GPTQ) 9.35 74.74 75.73 45.55 72.14 45.31 69.92 63.22 39.6 60.78 52.54 65.92 70.43 53.12 59.5			9.66														59.52
SERQ-MXFP4 9.79 75.17 74.81 46.88 69.57 44.28 68.85 63.85 40 60.43 49.8 61.92 64.9 51.32 56.1		SERO (GPTO)	9 35	74.74	75.73	45.55	72.14	45.31	69.92	63.22	39.6	60.78	52.54	65.92	70.43	53.12	59 55

Table 9: Accuracy results of distribution flattening methods when tested with W4A4 settings.

					0.61	4 C	C	D '	4					MMT I'		
								Reasoni	-					MMLU		
Model	Methods	PPL (↓)	BoolQ (†)	PIQA (†)	SIQA (†)	ARC-e (↑)	ARC-c	HellaS.	WinoG. (↑)	OBQA (†)	<b>Avg.</b> (↑)	Human.	Other (†)	SocialS.	STEM (†)	<b>Avg.</b> (↑)
	baseline	5.47	77.74	79.05	46.11	74.49	46.25	76	68.9	44.2	64.09	39.72	47.12	47.45	34.28	41.83
LLaMA-2-7B	SmoothQ SmoothQ(g128) QuaRot SpinQuant SERQ-MXFP4	1.51e+4 7.49 6.15 6 6.22	44.16 68.59 72.84 73.8 73	49.95 74.48 77.2 76 77.2	34.49 41.45 33.06 44.1 44.17	25.97 65.4 71.59 43.6 70.2	27.82 39.85 43 71.3 43.69	26.39 67.99 72.3 73.2 72.59	47.51 61.64 64.64 65.4 68.19	27.2 37.8 41.6 40.4 41	35.44 57.15 59.53 61 61.26	25.48 28.63 32.52 33.9 32.16	22.98 31.9 36.85 38.5 40.01	29.48 32.92 36.72 37.5 39.1	26.55 29.12 28.86 29.5 31.43	26.04 30.4 33.58 34.8 35.25
0-Shot Common Sense Reasoning tasks MMLU																
		PPL	PoolO	DIOA					WinoG.	OPOA	Ava	Uuman		SocialS.	CTEM	Avg.
Model	Methods	(\psi)	(†)	(†)	(†)	(†)	(†)	(†)	(†)	(†)	(↑)	(†)	(†)	(†)	(†)	(†)
	baseline	4.88	80.61	80.52	47.39	77.53	49.23	79.38	72.38	45.2	66.53	47.89	59.29	61.16	42.21	52.04
LLaMA-2-13B	SmoothQ SmoothQ(g128) QuaRot SpinQuant SERQ-MXFP4	1.32e+4 6.31 5.41 5.2 5.39	38.47 75.23 78.47 78.2 77.86	49.24 76.93 78.89 79.3 78.73	34.7 43.19 33.32 46.3 45.19	26.81 70.5 73.7 49 75.59	27.73 44.45 46.25 76.3 48.12	25.72 72.52 76.29 77.1 76.18	47.75 68.43 70.48 69.5 70.09	25.8 39 43 42.8 43	34.53 61.28 62.55 64.8 64.35	24.48 35.9 43.68 43.5 43.59	26.62 44.13 52.85 53.1 52.33	23.4 45.76 55.41 55.4 54.6	23.85 35.68 39.11 39.1 40.28	23.85 39.83 47.25 47.8 47.19
0-Shot Common Sense Reasoning tasks MMLU																
Model	Methods	PPL			SIQA	ARC-e	ARC-c	HellaS.	WinoG.				Other	SocialS.		Avg.
		(\psi)	(†)	(†)	(†)	(†)	(†)	(†)	(↑)	(†)	(†)	(↑)	(†)	(↑)	(†)	(†)
	baseline	6.13	81.07	80.74	47.08	77.69	52.99	79.2	73.48	45	67.16	54.81	70.55	73.25	53.89	62.13
LLaMA-3-8B	SmoothQ SmoothQ(g128) QuaRot SpinQuant SERQ-MXFP4	4.71e+5 17.26 8.41 8.26 7.63	50.61 61.83 70.49 73.4 76.15	51.69 64.2 77.04 75.2 77.2	33.78 39.92 32.96 44.4 44.11	24.2 46.59 69.57 72 72.69	26.02 30.38 43.26 46.9 45.39	26.65 55.61 72.22 71.9 75.13	49.57 59.59 64.64 67.7 68.43	28.2 33.6 42.8 42.4 42.6	36.34 48.97 59.12 61.75 62.71	26.65 29.33 42.76 45.8 47.72	24.36 30.16 53.56 56.5 60.7	24.21 29.61 53.56 57.2 63.18	24.71 28.1 41.74 42.5 46.84	25.17 29.3 47.29 49.93 53.48
					0-Sho	t Comm	on Sense	Reasoni	ng tasks					MMLU		
Model	Methods	PPL (↓)	BoolQ (↑)	PIQA (↑)	SIQA (†)	ARC-e (↑)	ARC-c (↑)	HellaS.	WinoG.	OBQA (†)	Avg. (†)	Human.	Other (†)	SocialS.	STEM (†)	Avg. (†)
	baseline	9.75	63.67	74.48	42.84	60.44	36.18	63.73	59.83	37.4	54.82	34.92	41.1	39.78	32.29	36.76
LLaMA-3.2-1B	SmoothQ SmoothQ(g128) QuaRot SpinQuant SERQ-MXFP4	1.53e+5 69.22 13.17 13.47 13.71	39.14 51.62 60.21 60.1 60.03	51.8 56.09 69.7 68.8 68.77	32.7 35.26 39.92 39 40.58	26.35 36.99 54.17 50.2 54.08	28.16 26.11 32.17 30.6 31.83	25.44 37.9 55.65 55.4 55.34	50.12 49.33 55.25 55.5 56.67	31 27 33.2 32.2 32.2	35.59 40.04 50.03 48.95 49.94	24.59 24.65 26.55 26.4 27.55	24.46 26.1 27.62 27.7 28.23	24.89 23.24 25.71 26 27.49	23.44 23.63 26.74 25.5 25.5	24.37 24.43 26.64 26.38 27.23
					0-Sho	t Comm	on Sense	Reasoni	ng tasks					MMLU		
Model	Methods	PPL (↓)	BoolQ (↑)	PIQA (↑)	SIQA (†)	ARC-e (↑)	ARC-c (↑)	HellaS.	WinoG.	OBQA	Avg. (†)	Human.	Other	SocialS.	STEM (†)	Avg. (†)
	baseline	7.81	72.97	77.53	46.93	71.59	46.25	73.49	69.53	43	62.66	48.78	63.08	62.59	44.72	54.06
LLaMA-3.2-3B	SmoothQ SmoothQ(g128) QuaRot SpinQuant SERQ-MXFP4	3.73e+4 53.33 9.73 10.15 8.74	40.46 53.7 66.15 68.5 68.56	51.52 62.51 73.34 73 73.83	33.32 37.41 43.45 43.2 45.24	25.34 42.89 59.81 62.8 61.83	26.11 27.99 37.29 38.9 40.02	26.38 49.04 68.12 67.6 67.16	51.46 53.2 60.85 63.1 64.33	31.2 26.6 37.2 38 39	35.72 44.17 55.76 56.88 57.5	23.91 27.38 41.23 39.9 38.68	23.53 27.81 51.05 47.1 47.89	23.33 27.62 50.76 47.7 45.82	22.61 26.42 37.93 36.4 34.44	23.41 27.31 44.75 42.42 41.33