

CHANNEL-WISE MIXED-PRECISION QUANTIZATION FOR LARGE LANGUAGE MODELS

Anonymous authors

Paper under double-blind review

ABSTRACT

Large Language Models (LLMs) have demonstrated remarkable success across a wide range of language tasks, but their deployment on edge devices remains challenging due to the substantial memory requirements imposed by their large parameter sizes. Weight-only quantization presents a promising solution to reduce the memory footprint of LLMs. However, existing approaches primarily focus on integer-bit quantization, limiting their adaptability to fractional-bit quantization tasks and preventing the full utilization of available storage space on devices. In this paper, we introduce Channel-Wise Mixed-Precision Quantization (CMPQ), a novel mixed-precision quantization method that allocates quantization precision in a channel-wise pattern based on activation distributions. By assigning different precision levels to different weight channels, CMPQ can adapt to any bit-width constraint. CMPQ employs a non-uniform quantization strategy and incorporates two outlier extraction techniques that collaboratively preserve the critical information, thereby minimizing the quantization loss. Experiments on different sizes of LLMs demonstrate that CMPQ not only enhances performance in integer-bit quantization tasks but also achieves significant performance gains with a modest increase in memory usage. CMPQ thus represents an adaptive and effective approach to LLM quantization, offering substantial benefits across diverse device capabilities.

1 INTRODUCTION

Large Language Models (LLMs), trained on massive text corpora and containing up to hundreds of billions of parameters, have demonstrated exceptional performance across a wide range of language tasks (Brown, 2020; Chowdhery et al., 2023; Du et al., 2022; Hoffmann et al., 2022; Thoppilan et al., 2022; Touvron et al., 2023a). However, deploying LLMs directly on devices for inference presents a significant challenge due to their enormous parameter sizes (Frantar & Alistarh, 2023; Xia et al., 2024; Xiao et al., 2024; Li et al., 2024b). For instance, GPT-3 (Brown, 2020), with 175 billion parameters, requires 350 GB of memory in FP16 precision, which far exceeds the capacity of the latest NVIDIA H100 GPU with 96 GB of memory, let alone the capabilities of edge devices.

Low-precision weight-only quantization (Park et al., 2024; Lee et al., 2024; Huang et al., 2024) has emerged as a promising solution to address this challenge by converting model weights from high bit-width representations (e.g., FP16) to lower bit-widths (e.g., 3-bit), significantly reducing the memory requirements for on-device LLM inference. In this work, we focus on Post-Training Quantization (PTQ) (Dettmers et al., 2022; Yuan et al., 2023; Shao et al., 2023; Liu et al., 2024), which quantizes the pre-trained models without the need for retraining – a process that is often costly and resource-intensive. Most existing PTQ methods concentrate on integer-bit quantization and employ a uniform low-bit-width representation across all layers (Lin et al., 2024; Chee et al., 2024; Frantar et al., 2022), as illustrated in Figure 1(a), yielding promising performance in low-bit quantization tasks. However, these methods are limited in their adaptability to devices with additional storage capacity. For example, they can only offer solutions constrained to integer-bit quantization (e.g., 2-bit) even if the device could support models with an average of, e.g., 2.2-bit precision, failing to utilize the extra storage space to further reduce quantization loss.

Mixed-precision quantization (Ma et al., 2023; Wang et al., 2019; Dong et al., 2019; Zhang et al., 2021) inherently supports fractional bit quantization by allowing model weights to be quantized at different precisions. However, few works have focused on mixed-precision quantization specifically

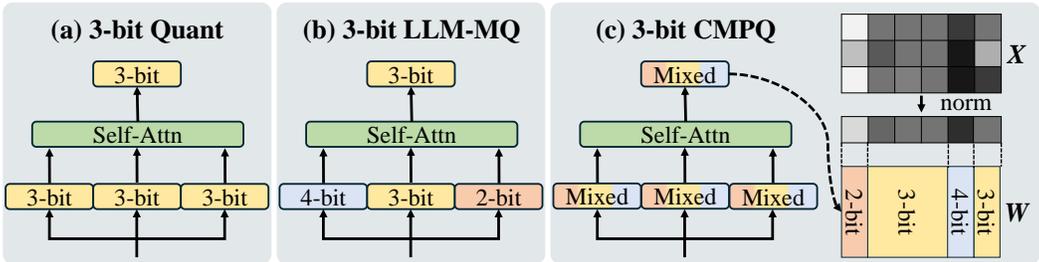


Figure 1: Illustration of different quantization approaches under a fixed bit-width constraint, such as 3 bits. (a) Standard quantization methods focus on algorithmic optimization to improve model performance, quantizing all layers uniformly to 3 bits. (b) LLM-MQ (Li et al., 2023) calculates layer-wise scores using first-order information and applies integer programming to assign lower bit-widths to less sensitive layers. (c) In contrast, our proposed CMPQ distributes the information loss evenly across layers by employing a channel-wise approach. This method assigns varying bit-widths within each layer based on activation distribution, ensuring that no single layer experiences significant information loss.

tailored for LLMs. LLM-MQ (Li et al., 2023) calculates the first-order information of layers at different bit-widths and uses integer programming to allocate layer-wise precision, as illustrated in Figure 1(b). Nevertheless, the gradient of a converged LLM is approximately zero, making it challenging for LLM-MQ to effectively differentiate the sensitivities of each layer. Furthermore, as demonstrated in Section 2, the additional quantization loss introduced by low-bit quantization can further degrade model performance.

To provide a mixed-precision method capable of adapting to any bit-width constraint without compromising the performance, we propose Channel-Wise Mixed-Precision Quantization (CMPQ). We first observe that different channels in the weight matrix have varying impacts on model performance, and that assigning higher precision to salient channels can enhance the performance of quantized LLMs. Inspired by this, CMPQ performs mixed-precision quantization on a channel-wise basis, as shown in Figure 1(c). Specifically, we compute the L_2 -norm of the activation for each layer and allocate high (or low) precision to channels with large (or small) activation norms. For the quantization process, we adopt a non-uniform quantization approach for each channel, accounting for the non-uniform nature of weight distributions. To further improve performance, we design two outlier extraction methods that separately focus on preserving activation-based outliers and quantization-aware outliers. Our main contributions are summarized as follows.

- We propose CMPQ, a mixed-precision quantization method designed to adaptively quantize LLMs to any specified bit-width. CMPQ performs channel-wise quantization, allocating precision based on activation distributions to optimize the performance.
- CMPQ incorporates a non-uniform quantization strategy along with two outlier protection methods that collaboratively preserve critical weights in high precision, thereby reducing quantization loss.
- We conduct extensive experiments to empirically validate the effectiveness of CMPQ, highlighting two key advantages: (1) channel-wise quantization based on sensitivity significantly improves performance in integer-bit quantization tasks, and (2) a modest increase in storage overhead results in substantial performance gains.

2 RELATED WORKS

Post-Training Quantization Quantization is a model compression technique that modifies the vector or matrix representations of a pre-trained model to improve inference efficiency. It can be broadly categorized into two workflows: Quantization-Aware Training (QAT) (Liu et al., 2023; Xu et al., 2024; Li et al., 2024c; Dettmers et al., 2024) and Post-Training Quantization (PTQ) (Yao et al., 2022; Xiao et al., 2023; Huang et al., 2024; Hooper et al., 2024). QAT involves retraining the model while adjusting its parameters during quantization, which is resource-intensive and often impractical for LLMs. In contrast, PTQ focuses on quantizing pre-trained models without the need for retraining, making it a more feasible approach for resource-constrained scenarios. However, the complexity of

LLMs necessitates specialized approaches for effective quantization (Zhou et al., 2024). We focus on weight-only PTQ methods (Lee et al., 2024; Park et al., 2024; Dettmers et al., 2023; Tseng et al., 2024). One of the early advancements in this domain is GPTQ (Frantar et al., 2022), an improvement over OBQ (Frantar & Alistarh, 2022), which determines the optimal quantization order per row of the weight matrix based on reconstruction error relative to the Hessian matrix of unquantized weights. QuIP (Chee et al., 2024) further refines GPTQ by introducing an optimal adaptive method for a quadratic proxy objective. This method enhances quantization effectiveness by ensuring incoherence between the weight and Hessian matrices, achieved through random orthogonal matrix multiplication. AWQ (Lin et al., 2024) addresses the varying importance of weight channels for performance by employing a reparameterization technique, selecting coefficients via grid search to minimize reconstruction errors efficiently. SqueezeLLM (Kim et al., 2024) approaches quantization by storing outliers in a full-precision sparse matrix while applying non-uniform quantization to the remaining weights. Although these methods demonstrate promising performance in low-bit quantization tasks, they are limited in their adaptability to devices with additional storage capacity. Specifically, they cannot be easily extended to fractional-bit tasks, which prevents them from leveraging extra storage to further reduce quantization loss and improve model performance.

Mixed-Precision Quantization Quantizing models to low precision uniformly can lead to significant accuracy degradation (Habi et al., 2020; Qu et al., 2020; Hu et al., 2021). To address this, Mixed-Precision Quantization (MPQ) has emerged as a promising approach, aiming to reduce model size and computational costs while maintaining or even improving accuracy by assigning different bit widths to weights (Wang et al., 2019; Li et al., 2024a; Rakka et al., 2022). Existing methods primarily focus on relatively small DNNs, such as MobileNet (Wang et al., 2019; Howard, 2017) and ResNet (Wu et al., 2018; He et al., 2016), and argue that MPQ is necessary for different layers, as each layer exhibits varying levels of redundancy and performs differently on hardware (Wang et al., 2019). To determine the appropriate per-layer precision, these methods rely on techniques like Neural Architecture Search (NAS) (Wu et al., 2018), periodic function regularization (Naumov et al., 2018), and second-order sensitivity analysis (Dong et al., 2019; 2020; Yao et al., 2021). However, these approaches are difficult to generalize to the quantization of LLMs due to the extensive computational resources and large search space they require (Gholami et al., 2022). While protecting outliers in FP16 can be considered a form of mixed-precision, few studies have applied mixed-precision strategies to the majority of weights. LLM-MQ (Li et al., 2023) uses integer programming to allocate layer-wise precision based on first-order information. However, as discussed in Section 3.1.2, it struggles to capture the varying sensitivities of individual layers. To overcome the limitations of large search spaces and performance degradation, we propose a channel-wise mixed-precision strategy that protects salient channels within each layer, ensuring the preservation of critical information.

3 METHOD

In this section, we begin with preliminary studies to investigate the impact of channel-wise mixed-precision quantization compared to layer-wise mixed-precision quantization (Section 3.1.2). The empirical results align with our intuition: channel-wise mixed-precision quantization more effectively harnesses the potential of mixed-precision. Subsequently, in Section 3.2, we provide a detailed introduction to our proposed Channel-Wise Mixed-Precision Quantization method.

3.1 PRELIMINARY

3.1.1 N -BIT QUANTIZATION

Quantization typically involves mapping a continuous set of values \mathbf{W} from a higher bit-width (e.g., 16-bit floating point) to a discrete set of values $Q(\mathbf{W})$ at a lower bit-width (e.g., 4-bit integers). The most commonly used uniform quantization (Krishnamoorthi, 2018) can be expressed as:

$$Q(\mathbf{W}) = \left\lfloor \frac{\mathbf{W}_{\text{FP16}} - \min(\mathbf{W}_{\text{FP16}})}{\Delta} \right\rfloor,$$

where Δ is the quantization step size, determined by the range of the original values \mathbf{W} and the desired bit-width N . Specifically, Δ can be computed as:

$$\Delta = \frac{\max(\mathbf{W}_{\text{FP16}}) - \min(\mathbf{W}_{\text{FP16}})}{2^{N-1} - 1}.$$

However, LLMs typically exhibit non-uniform weight distributions (Kim et al., 2024; Dettmers et al., 2024). In such cases, the presence of large magnitude values can lead to inefficient quantization, where certain bins or bit combinations are underutilized, resulting in suboptimal representation with few or no values assigned to some bins. To address this issue and better preserve the original weight information, we adopt non-uniform quantization. For a given vector \mathbf{x} , we compute a 2^N -bit quantized representation, denoted as $\mathbf{q} = \{q_1, \dots, q_{2^N}\}$, and obtain its lower-precision approximation through the following procedure:

$$Q(\mathbf{x}) = (Q(x_i)), \text{ where } Q(x_i) = \min_{q_j \in \mathbf{q}} |x_i - q_j|.$$

3.1.2 PRELIMINARY STUDY OF MIXED-PRECISION QUANTIZATION

To explore the effect of mixed-precision quantization, we conduct preliminary experiments on two OPT models (Zhang et al., 2022) and compare the perplexity evaluation. We focus on 3-bit quantization, where precision is selected from $\{2, 3, 4\}$ -bit.

In Li et al. (2023), each linear layer is associated with information loss scores under different precisions, and their approach models the average bit width as a constraint in an integer programming problem. We implement this strategy with the non-uniform quantization, and the performance is reported in Table 1 as **w/ IntProg**. Additionally, Lin et al. (2024) observe that retaining salient weights based on activation distributions in higher precision significantly enhances quantized performance. For each layer, we compute the channel-wise L_2 -norm of activations and select the top (and bottom) $k\%$ of channels. These channels are quantized to 4-bit (high precision) or 2-bit (low precision) respectively. Results for $k = 1$ and $k = 10$ are presented in Table 1.

Table 1: Preliminary study of various mixed-precision 3-bit quantization methods. **w/ IntProg** denotes quantization using an integer programming solution for each layer. **w/ 10% 2-bit, 10% 4-bit** and **w/ 1% 2-bit, 1% 4-bit** indicate quantization where 10% (1%) of channels are 4-bit and 10% (1%) are 2-bit, based on the activation distribution. Best performances are in bold, with underlined text showing the second best.

Method	OPT-2.7		OPT-6.7	
	Wiki (\downarrow)	C4 (\downarrow)	Wiki (\downarrow)	C4 (\downarrow)
3-bit	13.45	14.08	11.48	12.28
w/ IntProg	14.75	14.88	12.69	13.13
w/ 10% 2-bit, 10% 4-bit	13.53	14.27	11.61	12.42
w/ 1% 2-bit, 1% 4-bit	13.38	14.05	11.45	12.26

From Table 1, we observe that applying integer programming and assigning a fixed bit-width to each layer is insufficient for achieving effective mixed-precision quantization, and may even perform worse than using 3-bit quantization across all layers. This could be attributed to two factors: (i) the per-layer scores defined in Li et al. (2023) struggle to effectively distinguish the sensitivity of different layers, as the gradients in a converged LLM are nearly zero, and (ii) the additional information loss introduced by 2-bit quantization outweighs the compensatory gains from 4-bit quantization when compared to 3-bit quantization (Chee et al., 2024). On the other hand, quantizing each layer to different precisions based on activation distributions can yield better results than consistently using 3-bit across all layers. However, extending mixed-precision quantization to more channels could degrade performance, due to the same issue outlined in (ii).

3.2 CHANNEL-WISE MIXED-PRECISION QUANTIZATION

Our primary objective is to develop an algorithm that effectively utilizes mixed-precision quantization to adaptively compress LLMs under any given average bit constraint, including fractional bit-widths. Additionally, we still aim to achieve strong performance on integer-bit quantization tasks, such as 3-bit quantization, compared with existing works. To this end, we propose Channel-Wise Mixed-Precision Quantization (CMPQ). In this section, we first introduce the channel-wise non-uniform quantization method, followed by a detailed explanation of our outlier protection strategy.

3.2.1 CHANNEL-WISE NON-UNIFORM QUANTIZATION

From the observations in Section 3.1.2, we can draw two key intuitions: ❶ channel-wise mixed-precision quantization can enhance model performance compared to layer-wise mixed-precision quantization, and ❷ when implementing channel-wise quantization, it is advisable to limit the number of 2-bit channels to minimize the information loss. Based on these, we propose channel-wise non-uniform quantization.

Research has shown that the weight distributions in LLM layers exhibit non-uniform patterns. Previous approaches have primarily focused on uniform quantization (Chee et al., 2024; Frantar et al., 2022), which divides the weight range into evenly spaced bins. However, this approach is suboptimal, as it fails to account for the non-uniform nature of the weight distributions, and struggles to improve end-to-end latency in memory-bound LLM inference (Kim et al., 2024). Following Kim et al. (2024), we adopt non-uniform quantization. Specifically, for each channel $W_{i,:}$ in the weight matrix $W \in \mathbb{R}^{d_{in} \times d_{out}}$, we apply a K -means clustering algorithm, where the value of K is determined by the precision assigned to the channel (e.g., $K = 8$ for 3-bit quantization). After clustering, each weight in W is represented by its nearest centroid from the set of K centroids $\{q_1, \dots, q_K\}$.

Algorithm 1 Channel-wise Precision Allocation

Input: Activation norm vector $\mathbf{a} \in \mathbb{R}^{d_{in}}$, average bit-width constraint $b \in [2, 4]$

Output: Channel precision allocation vector $\mathbf{c} \in \mathbb{R}^{d_{in}}$

Initialize $\mathbf{c} = 3 \cdot \mathbf{1}$

if $b > 3$ **then**

 Compute the q -th quantile l of \mathbf{a} , where $q = 1 - (b - 3)$

 Set $c[\mathbf{a} > l] = 4$

else if $b < 3$ **then**

 Compute the q -th quantile s of \mathbf{a} , where $q = 3 - b$

 Set $c[\mathbf{a} < s] = 2$

else

 Compute the 1st quantile s and the 99th quantile l of \mathbf{a}

 Set $c[\mathbf{a} < s] = 2$ and $c[\mathbf{a} > l] = 4$

end if

We draw inspiration from Lin et al. (2024) and our preliminary studies to determine the precision for each channel through a simple yet effective approach. We sample a calibration set to perform forward propagation on the LLMs, obtaining the activation matrix $X \in \mathbb{R}^{n \times d_{in}}$ for each layer weight matrix W . The per-channel L_2 -norm of X is then computed, yielding a 1-dimensional vector $\mathbf{a} \in \mathbb{R}^{d_{in}}$. Based on this, we calculate the channel precision allocation vector \mathbf{c} as described in Algorithm 1, and apply non-uniform quantization to quantize each channel $W_{i,:}$ to c_i bits. The key idea is that when the average bit-width b exceeds 3, we focus on quantizing the most salient channels – determined by the activation distribution – into higher precision to enhance model performance. Conversely, when the average bit-width b is below 3, we concentrate on quantizing less critical channels into lower precision to minimize quantization loss. For 3-bit quantization, we protect approximately 1% of the salient weight channels by assigning them 4-bit precision for improved performance, as motivated by the findings in Table 1.

It is worth noting that we intentionally avoid protecting additional channels with 4-bit precision and avoid incorporating equivalent channels in 2-bit quantization for compensation due to the following reasons: (i) Only a small fraction of weights are salient (Lin et al., 2024), and protecting them would significantly reduce quantization loss. (ii) The additional information loss incurred by 2-bit quantization outweighs the compensatory benefits gained from 4-bit quantization.

3.2.2 OUTLIER PROTECTION

Another key challenge in low-bit LLM quantization is the protection of outlier values (Bondarenko et al., 2021; Wei et al., 2022; Dettmers et al., 2022). Previous studies have demonstrated that naively quantizing weights with a large dynamic range significantly degrades performance, particularly at low precisions (Kim et al., 2024). However, in some cases, retaining a small fraction (less than 1%) of

270 outlier weights in FP16 has been shown to reduce up to 75% of the total quantization error (Dettmers
 271 et al., 2023). This suggests that extracting outliers prior to quantization can mitigate their negative
 272 impact and minimize quantization loss. Consistent with prior works (Li et al., 2023; Kim et al., 2024),
 273 we retain 0.5% of outliers in high precision (16-bit), while applying quantization to the remaining
 274 weights. Our approach focuses on protecting two types of outliers: activation-based outliers and
 275 quantization-based outliers.

276 As illustrated in Figure 2, we observe that outliers exhibit
 277 a channel-wise pattern – if an outlier appears in a channel,
 278 it consistently occurs across all tokens. Table 1 empirically
 279 demonstrates that preserving salient weights based on the
 280 activation distribution helps mitigate quantization loss.
 281 Motivated by these findings, we introduce an activation-
 282 based outlier detection method, which identifies outliers
 283 O_{act} from the weight matrix W . Specifically, we select
 284 channels corresponding to the top 0.45% largest values
 285 in the activation’s L2-norm vector α , and preserve these
 286 channels in FP16 precision.

287 In addition to selecting channel-wise outliers, we also
 288 investigate the protection of a small subset of quantization-
 289 sensitive outliers. Though we introduced our non-uniform
 290 quantization method, a small fraction of weights exhibit
 291 significantly larger magnitudes compared to the majority.
 292 These large weights can distort the clustering process by
 293 shifting centroids away from the bulk of the weight distri-
 294 bution, thereby negatively impacting the performance of
 295 the quantized LLM. A conventional approach to outlier protection involves the removal of weights
 296 based solely on their magnitude. However, instead of simply eliminating high-magnitude outliers,
 297 our objective is to identify and remove those that most adversely affect the quantization process.

298 To achieve this, we apply another K -means clustering step prior to quantization. Specifically,
 299 given $W' = W - O_{\text{act}}$ that represents the remaining weights after activation outlier removal, we
 300 use Algorithm 1 to determine the channel precisions c . We then apply channel-wise non-uniform
 301 quantization based on c to obtain the quantized model W'_q . We identify the set of outliers O_q in W'
 302 corresponding to the top 0.05% of the largest values in $|W' - W'_q|$. This approach preserves these
 303 magnitude-based outliers in FP16 format, not only to mitigate their influence on model output but
 304 also to ensure that the centroids $\{q_1, \dots, q_K\}$ better represent the majority of the weights, rather than
 305 being skewed by a small number of outliers. Two types of outliers $O = O_{\text{act}} + O_q$ are removed from
 306 the weight matrix W , and the remaining weights then undergo the quantization process to obtain W_q .
 307 $W_q + O$ is used for the final inference. Notably, the overhead associated with this decomposition is
 308 minimal, as the number of outlier values is relatively small, typically around 0.5% of the total values.

3.2.3 DISCUSSION

309 Our method only requires forward propagation and does not depend on backpropagation, which is
 310 necessary for many existing quantization techniques (Li et al., 2023; Kim et al., 2024). Consequently,
 311 the memory requirements for our proposed CMPQ during quantization are moderate; for instance,
 312 loading the OPT-6.7B model necessitates 12.4 GB of memory, whereas the backward pass for
 313 the same model requires 49.61 GB of memory. Additionally, CMPQ has minimal reliance on the
 314 calibration set, as it only measures the L_2 -norm per channel, thereby mitigating the risk of overfitting.
 315 For a comparison with backpropagation-dependent methods, refer to Section 4.4, and for an analysis
 316 of CMPQ’s robustness with respect to the calibration dataset, see Section 4.5.

4 EXPERIMENTS

4.1 EXPERIMENT SETUP

322 **LLM Models and Datasets.** We perform our experiments on two models from the OPT family (Zhang
 323 et al., 2022) (OPT-2.7B and OPT-6.7B) and two models from the LLaMA2 family (Touvron et al.,

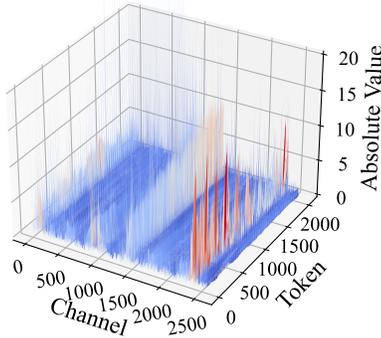


Figure 2: Magnitude of absolute activation values of the self_attn.out_proj layer in the third layer of OPT-2.7B.

2023b) (LLaMA2-7B and LLaMA2-13B). The evaluation of the quantized models is based on perplexity across two language generation tasks, WikiText-2 (Merity et al., 2016) and C4 (Raffel et al., 2020), as perplexity is a widely recognized metric for assessing the LLM performance. For calibration, we follow previous works (Chee et al., 2024; Frantar et al., 2022) and use a set of 128 randomly selected 2048-token segments from the C4 dataset, which contains generic text data from web crawls. This ensures consistency in comparison with baselines and avoids the use of task-specific data when quantizing other datasets. All experiments are implemented in PyTorch (Paszke et al., 2019) and executed on two A6000 GPUs, with performance monitoring handled by the Torch CUDA profiler. We extend our evaluation in Appendix A.3 to include quantization results for other OPT models, scaling up to 30B parameters, as well as a newer model, LLaMA3-8B (AI@Meta, 2024).

Baselines. We evaluate the proposed CMPQ against several post-training quantization methods that do not rely on backpropagation, including Round-to-Nearest (RTN), GPTQ (Frantar et al., 2022), AWQ (Lin et al., 2024), and QuIP (Chee et al., 2024). Since these methods are specifically designed for integer bit-width quantization, we restrict the comparison to {2, 3, 4}-bit settings. Additionally, we compare CMPQ with a mixed-precision quantization method tailored for LLMs, LLM-MQ (Li et al., 2023), focusing on performance in fractional bit-width quantization. In Section 4.4, we extend the comparison to also include SqueezeLLM (Kim et al., 2024), a state-of-the-art gradient-based method, and discuss the trade-offs between memory cost and quantization performance.

4.2 MAIN RESULTS

Table 2 presents the main results comparing CMPQ with post-training quantization baselines. Overall, our proposed CMPQ consistently outperforms all baselines, particularly in the 2-bit quantization tasks. Notably, while QuIP is specifically designed for low-bit quantization, it performs poorly on the WikiText-2 task using 2-bit quantized small LLM models (OPT-2.7B). In contrast, CMPQ achieves significantly better results on this task, despite using the same calibration dataset from C4. This highlights that CMPQ is less sensitive to the choice of the calibration set, as it relies solely on measuring the activation per-channel L_2 -norm, which generalizes more effectively across different dataset distributions. Furthermore, although LLM-MQ is designed for mixed-precision quantization, it struggles to achieve competitive performance in 3-bit quantization compared to other baselines. This limitation arises from its precision allocation strategy, which is based exclusively on first-order information that is hard to distinguish layer sensitivities of a converged LLM.

4.3 NON-INTEGGER QUANTIZATION

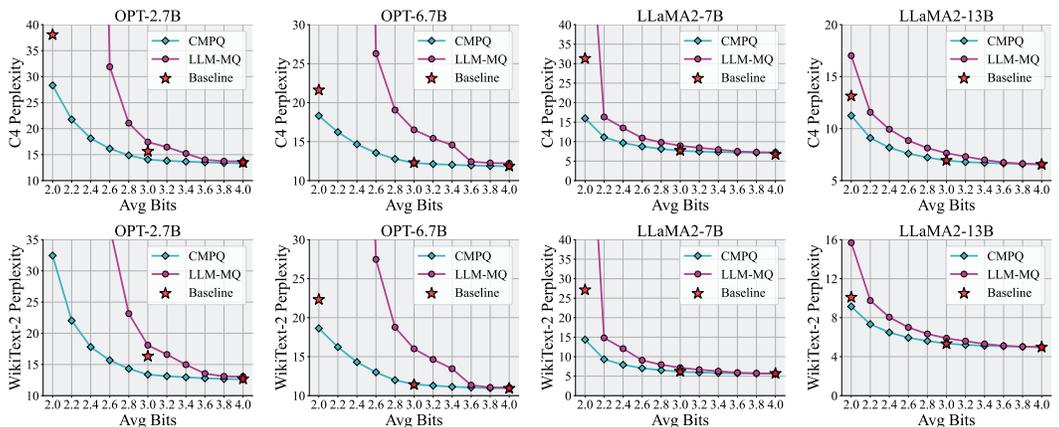


Figure 3: Comparison of C4 perplexity between CMPQ and LLM-MQ for fractional bit-width quantization. The red star indicates the best performance achieved by integer-based baselines at {2, 3, and 4} bits.

In Figure 3, we compare the perplexity of LLMs quantized by CMPQ and LLM-MQ under fractional bit-width constraints. First, it is evident that CMPQ consistently outperforms LLM-MQ across various bit-widths and models. LLM-MQ quantizes entire layers to a fixed precision, which can result in significant information loss within a single layer, negatively impacting the overall model

Table 2: Perplexity (\downarrow) comparison of OPT and LLaMA2 models quantized to {2, 3, 4}-bit precision using various quantization methods on the C4 and WikiText-2 datasets. Bold font indicates the best performance across all methods, while underlined results denote the second-best.

Method	Avg. Bit	OPT-2.7B		OPT-6.7B		LLaMA2-7B		LLaMA2-13B	
		Wiki	C4	Wiki	C4	Wiki	C4	Wiki	C4
FP16	16	12.47	13.17	10.86	11.74	5.47	6.97	4.88	6.47
RTN	2	>100	>100	>100	>100	>100	>100	>100	>100
GPTQ	2	>100	>100	>100	>100	36.77	35.7	13.67	16.45
AWQ	2	–	–	–	–	>100	>100	>100	>100
QuIP	2	>100	<u>38.07</u>	<u>22.33</u>	<u>21.62</u>	<u>27.12</u>	<u>31.33</u>	<u>10.09</u>	<u>13.13</u>
LLM-MQ	2	>100	>100	>100	>100	96.61	85.16	15.69	17.02
CMPQ	2	32.46	28.32	18.63	18.31	14.37	15.97	9.14	11.25
RTN	3	>100	>100	>100	>100	>100	>100	10.68	12.50
GPTQ	3	17.09	18.14	14.87	17.13	6.25	7.97	6.17	7.06
AWQ	3	<u>16.32</u>	16.28	11.41	<u>12.30</u>	<u>6.24</u>	7.84	5.32	6.94
QuIP	3	17.44	<u>15.63</u>	11.51	13.30	6.80	<u>7.75</u>	5.65	7.25
LLM-MQ	3	18.09	17.42	16.01	16.51	7.16	8.94	5.89	7.64
CMPQ	3	13.38	14.05	<u>11.45</u>	12.26	6.14	7.66	<u>5.34</u>	6.93
RTN	4	16.69	18.75	12.15	14.40	6.12	7.72	5.20	6.83
GPTQ	4	12.93	14.99	11.49	13.16	5.72	7.23	5.08	6.74
AWQ	4	12.73	<u>13.48</u>	10.93	<u>11.86</u>	5.72	7.13	4.98	<u>6.56</u>
QuIP	4	<u>12.69</u>	14.55	10.98	12.86	5.72	6.69	5.29	6.83
LLM-MQ	4	13.06	13.70	11.04	12.22	<u>5.68</u>	7.22	4.98	6.58
CMPQ	4	12.63	13.33	<u>10.95</u>	11.83	5.61	<u>7.10</u>	4.98	6.55

performance. In contrast, CMPQ allocates mixed-precision in a channel-wise manner, distributing the information loss more evenly across layers and avoiding a substantial loss in any single layer. Another key observation, visible when examining the transition from 2-bit to 2.2-bit quantization, highlights the advantage of mixed-precision. Introducing just a 10% increase in storage overhead at lower bit-widths can lead to significant performance gains. For instance, in the case of LLaMA2-7B, CMPQ yields a 30% improvement in perplexity (from 15.97 to 11.11), while LLM-MQ also shows a dramatic improvement, moving from poor performance (85.16) to performance that even surpasses the baseline (16.32). This demonstrates that mixed-precision quantization can trade a small increase in storage overhead for a substantial boost in performance – something that is not achievable with integer-only bit quantization methods.

As the average bit-width increases, the performance of both methods converges and approaches that of the baseline at 4-bit quantization. This indicates that quantization techniques face greater challenges at lower bit-widths. The strong performance of CMPQ at 2-bit and 3-bit quantization demonstrates its effectiveness, particularly in scenarios where lower bit-widths are required.

4.4 COMPARISON WITH SQUEEZELLM

While we primarily compare with baselines that do not rely on backpropagation, we acknowledge that gradient information, at the cost of extra resources, can indeed enhance the performance of quantized LLMs. In Table 3, we compare our method with the state-of-the-art baseline, SqueezeLLM (Kim et al., 2024), which also employs non-uniform quantization but uses gradient information to weight the clustering process, safeguarding more sensitive weights. Additionally, we report the memory requirements for loading LLMs and performing backpropagation in FP16 precision¹.

As shown in Table 3, SqueezeLLM outperforms CMPQ across various models at different quantization levels. However, this improvement comes at a significant cost: SqueezeLLM requires four times the memory for the quantization process, making it impractical for larger models, especially under

¹The memory requirements were calculated using the Hugging Face platform (<https://huggingface.co/spaces/hf-accelerate/model-memory-usage>).

Table 3: Comparison of WikiText-2 perplexity between SqueezeLLM (abbreviated as SqzLLM) and CMPQ across various average bit-widths and models. The Memory column includes the total model size and memory requirements for backpropagation with a batch size of 1. For CMPQ, we also report performance at 2.2-bit and 3.2-bit to facilitate trade-off discussions. Bold font indicates the best performance across all methods, while underlined results denote the second-best.

Models	Memory (GB)	SqzLLM	CMPQ	SqzLLM	CMPQ	SqzLLM	CMPQ
		2	2/2.2	3	3/3.2	4	4
OPT-2.7B	4.94/19.76	–	<u>32.46/22.03</u>	13.43	<u>13.38/13.12</u>	12.60	12.63
OPT-6.7B	12.40/49.61	–	<u>18.63/16.23</u>	<u>11.31</u>	<u>11.45/11.29</u>	10.92	10.95
LLaMA2-7B	12.37/49.48	<u>10.79</u>	14.37/ 9.32	<u>5.96</u>	6.14/ 5.95	5.57	5.61
LLaMA2-13B	24.02/96.07	<u>7.91</u>	9.14/ 7.33	<u>5.23</u>	5.34/ 5.21	4.96	4.98

resource constraints. In contrast, CMPQ offers a more efficient solution when memory is limited, requiring only 1/4 of the computational resources. Moreover, if an additional 10% of storage space is available, CMPQ can achieve better performance than SqueezeLLM, particularly in low-memory environments. At higher precision, such as 4-bit quantization, the tradeoff between computational resource requirements and model storage is less pronounced. The maximum performance gain of SqueezeLLM over CMPQ is marginal – only $(5.61 - 5.57)/5.61 = 0.71\%$. This minimal improvement renders the substantial additional resource demands of SqueezeLLM unnecessary.

Overall, while CMPQ may not surpass SqueezeLLM in the integer-only quantization tasks, the method offers significant advantages by eliminating the 300% increase in computational overhead associated with SqueezeLLM’s gradient-based approach, with only a 10% increase in storage. Furthermore, CMPQ is more versatile, as it applies to non-integer bit quantization tasks, making it a more practical option for a wider range of scenarios.

4.5 DATA EFFICIENCY AND GENERALIZATION

Data Efficiency for the Calibration Set. In Table 4, we present a data efficiency analysis based on the number of data samples in the calibration datasets and compare the perplexity of the LLaMA2-7B model under 3-bit quantization. Although a calibration set of 128 data samples is used consistently throughout the paper, our method typically achieves the desired quantization performance with as few as single-digit sample sizes. This efficiency stems from the fact that we do not rely on regression or backpropagation; instead, we only measure the activation norm from the calibration set, making the process highly data-efficient. In contrast, both GPTQ and AWQ require more than 50 data points for calibration, as reported in (Kim et al., 2024).

Table 4: Data efficiency analysis of calibration datasets: Comparing CMPQ and SqueezeLLM for perplexity on C4 and WikiText-2 with 3-bit quantization of the LLaMA2-7B across varying calibration set sizes.

Methods	Tasks	Nmuber of calibration samples					
		1	2	5	10	20	100
SqzLLM	Wiki	6.41	6.22	6.20	6.16.7	6.16	6.18
	C4	7.89	7.81	7.73	7.72	7.72	7.72
CMPQ	Wiki	6.18	6.14	6.16	6.14	6.14	6.14
	C4	7.71	7.65	7.66	7.65	7.65	7.66

Robustness to the Calibration Set. We evaluate the robustness of CMPQ by analyzing its performance using different calibration sets. Specifically, we compare CMPQ with QuIP in quantizing two OPT models into 2-bit representations. The results, presented in Table 5, demonstrate that CMPQ consistently outperforms QuIP in low-bit quantization across different calibration sets.

As the size of LLMs increases, both methods demonstrate improved robustness. However, a notable difference emerges in their performance under varying conditions. While QuIP performs effectively on 2-bit quantization when the evaluation dataset aligns with the calibration set (see Table 2), it experiences a significant performance decline, and may even diverge, when tested on a different dataset. This decline can be attributed to QuIP’s dependence on the Hessian matrix derived from

the calibration set, which makes it highly sensitive to changes in the dataset. In contrast, CMPQ exhibits greater resilience, relying only on the average activation L_2 -norm from the calibration set – a measure that generalizes more robustly across different datasets.

Table 5: Robustness analysis of the calibration dataset. We compare CMPQ with QuIP for 2-bit quantization tasks and report the perplexity differences across two different calibration datasets.

Models	Eval Calib	QuIP		CMPQ	
		Wiki	C4	Wiki	C4
OPT-2.7B	Wiki	32.84	242.69 (+204.62)	29.62	27.56 (-0.76)
	C4	>1000(+1000)	38.07	32.46(+2.84)	28.32
OPT-6.7B	Wiki	22.16	107.56 (+85.94)	18.86	18.39 (+0.08)
	C4	22.33(+0.17)	21.62	18.63 (-0.23)	18.31

4.6 ABLATION STUDY OF OUTLIER PROTECTION

In this section, we conduct experiments on the OPT-2.7B and OPT-6.7B models to analyze the impact of two distinct types of outliers on quantization performance. For a consistent comparison, when we remove one type of outlier, we retain the other type, ensuring that it constitutes 0.5% of the entire weight matrix. The results, presented in Table 6, demonstrate that both types of outliers contribute to enhancing the performance of quantized LLMs, particularly in low-bit settings. Notably, protecting both types of outliers yields the best results in general. This is because each type of outlier addresses different aspects: activation-based outliers safeguard salient weights, while quantization-based outliers ensure that the clustering process during quantization is not distorted by extreme values, thereby focusing on the majority of weights. In summary, these two outlier protection strategies complement each other, working in tandem to improve the overall model performance.

Table 6: Ablation of the outlier protection strategies. Best performances are in bold, with underlined text showing the second best.

Method	Avg. Bit	OPT-2.7B		OPT-6.7B	
		Wiki	C4	Wiki	C4
w/o Outlier	3	13.84	14.47	11.58	12.49
w/o O_q	3	13.59	14.28	<u>11.46</u>	12.36
w/o O_{act}	3	<u>13.39</u>	<u>14.05</u>	11.49	12.23
CMPQ	3	13.38	14.05	11.45	<u>12.26</u>
w/o Outlier	4	12.87	13.44	11.18	11.94
w/o O_q	4	<u>12.68</u>	<u>13.34</u>	<u>10.96</u>	<u>11.85</u>
w/o O_{act}	4	12.69	13.38	11.06	11.86
CMPQ	4	12.63	13.33	10.95	11.83

5 CONCLUSION AND FUTURE WORK

In this work, we focused on mixed-precision quantization and aimed to design an algorithm capable of adapting to any bit-width constraint. We observed that different weight channels had varying impacts on model performance, and that activation distributions helped identify salient channels. Building on these insights, we proposed CMPQ, which integrated a channel-wise non-uniform quantization strategy. To further enhance performance, CMPQ introduced two types of outliers that collaboratively preserved critical information. Experimental results showed that CMPQ harnessed the potential of mixed-precision quantization in two key ways: (1) it achieved superior performance in integer-bit quantization tasks, and (2) it delivered significant performance improvements with only a modest increase in memory requirements. In the future, our focus will be on deploying CMPQ-quantized LLMs on real-world devices. This will involve addressing key challenges such as engineering hardware acceleration and designing efficient lookup table kernels to optimize performance.

REFERENCES

- 540
541
542 AI@Meta. Llama 3 model card. 2024. URL [https://github.com/meta-llama/llama3/](https://github.com/meta-llama/llama3/blob/main/MODEL_CARD.md)
543 [blob/main/MODEL_CARD.md](https://github.com/meta-llama/llama3/blob/main/MODEL_CARD.md).
- 544 Yelysei Bondarenko, Markus Nagel, and Tijmen Blankevoort. Understanding and overcoming the
545 challenges of efficient transformer quantization. *arXiv preprint arXiv:2109.12948*, 2021.
- 546
547 Tom B Brown. Language models are few-shot learners. *arXiv preprint arXiv:2005.14165*, 2020.
- 548
549 Jerry Chee, Yaohui Cai, Volodymyr Kuleshov, and Christopher M De Sa. Quip: 2-bit quantization of
550 large language models with guarantees. *Advances in Neural Information Processing Systems*, 36,
551 2024.
- 552 Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam
553 Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, et al. Palm:
554 Scaling language modeling with pathways. *Journal of Machine Learning Research*, 24(240):1–113,
555 2023.
- 556 Tim Dettmers, Mike Lewis, Younes Belkada, and Luke Zettlemoyer. Gpt3. int8 (): 8-bit matrix
557 multiplication for transformers at scale. *Advances in Neural Information Processing Systems*, 35:
558 30318–30332, 2022.
- 559
560 Tim Dettmers, Ruslan Svirschevski, Vage Egiazarian, Denis Kuznedelev, Elias Frantar, Saleh Ashk-
561 boos, Alexander Borzunov, Torsten Hoefler, and Dan Alistarh. Spqr: A sparse-quantized represen-
562 tation for near-lossless llm weight compression. *arXiv preprint arXiv:2306.03078*, 2023.
- 563
564 Tim Dettmers, Artidoro Pagnoni, Ari Holtzman, and Luke Zettlemoyer. Qlora: Efficient finetuning
565 of quantized llms. *Advances in Neural Information Processing Systems*, 36, 2024.
- 566
567 Zhen Dong, Zhewei Yao, Amir Gholami, Michael W Mahoney, and Kurt Keutzer. Hawq: Hessian
568 aware quantization of neural networks with mixed-precision. In *Proceedings of the IEEE/CVF*
569 *international conference on computer vision*, pp. 293–302, 2019.
- 570
571 Zhen Dong, Zhewei Yao, Daiyaan Arfeen, Amir Gholami, Michael W Mahoney, and Kurt Keutzer.
572 Hawq-v2: Hessian aware trace-weighted quantization of neural networks. *Advances in neural*
573 *information processing systems*, 33:18518–18529, 2020.
- 574
575 Nan Du, Yanping Huang, Andrew M Dai, Simon Tong, Dmitry Lepikhin, Yuanzhong Xu, Maxim
576 Krikun, Yanqi Zhou, Adams Wei Yu, Orhan Firat, et al. Glam: Efficient scaling of language
577 models with mixture-of-experts. In *International Conference on Machine Learning*, pp. 5547–5569.
578 PMLR, 2022.
- 579
580 Elias Frantar and Dan Alistarh. Optimal brain compression: A framework for accurate post-training
581 quantization and pruning. *Advances in Neural Information Processing Systems*, 35:4475–4488,
582 2022.
- 583
584 Elias Frantar and Dan Alistarh. Sparsegpt: Massive language models can be accurately pruned in
585 one-shot. In *International Conference on Machine Learning*, pp. 10323–10337. PMLR, 2023.
- 586
587 Elias Frantar, Saleh Ashkboos, Torsten Hoefler, and Dan Alistarh. Gptq: Accurate post-training
588 quantization for generative pre-trained transformers. *arXiv preprint arXiv:2210.17323*, 2022.
- 589
590 Amir Gholami, Sehoon Kim, Zhen Dong, Zhewei Yao, Michael W Mahoney, and Kurt Keutzer. A
591 survey of quantization methods for efficient neural network inference. In *Low-Power Computer*
592 *Vision*, pp. 291–326. Chapman and Hall/CRC, 2022.
- 593
594 Hai Victor Habi, Roy H Jennings, and Arnon Netzer. Hmq: Hardware friendly mixed precision
595 quantization block for cnns. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow,*
596 *UK, August 23–28, 2020, Proceedings, Part XXVI 16*, pp. 448–463. Springer, 2020.
- 597
598 Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image
599 recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*,
600 pp. 770–778, 2016.

- 594 Jordan Hoffmann, Sebastian Borgeaud, Arthur Mensch, Elena Buchatskaya, Trevor Cai, Eliza
595 Rutherford, Diego de Las Casas, Lisa Anne Hendricks, Johannes Welbl, Aidan Clark, et al.
596 Training compute-optimal large language models. *arXiv preprint arXiv:2203.15556*, 2022.
597
- 598 Coleman Hooper, Sehoon Kim, Hiva Mohammadzadeh, Michael W Mahoney, Yakun Sophia Shao,
599 Kurt Keutzer, and Amir Gholami. Kvquant: Towards 10 million context length llm inference with
600 kv cache quantization. *arXiv preprint arXiv:2401.18079*, 2024.
- 601 Andrew G Howard. Mobilenets: Efficient convolutional neural networks for mobile vision applica-
602 tions. *arXiv preprint arXiv:1704.04861*, 2017.
603
- 604 Peng Hu, Xi Peng, Hongyuan Zhu, Mohamed M Sabry Aly, and Jie Lin. Opq: Compressing deep
605 neural networks with one-shot pruning-quantization. In *Proceedings of the AAAI conference on*
606 *artificial intelligence*, volume 35, pp. 7780–7788, 2021.
- 607 Wei Huang, Yangdong Liu, Haotong Qin, Ying Li, Shiming Zhang, Xianglong Liu, Michele Magno,
608 and Xiaojuan Qi. Billm: Pushing the limit of post-training quantization for llms. *arXiv preprint*
609 *arXiv:2402.04291*, 2024.
- 610 Sehoon Kim, Coleman Richard Charles Hooper, Amir Gholami, Zhen Dong, Xiuyu Li, Sheng Shen,
611 Michael W Mahoney, and Kurt Keutzer. Squeezellm: Dense-and-sparse quantization. In *Forty-first*
612 *International Conference on Machine Learning*, 2024.
613
- 614 Raghuraman Krishnamoorthi. Quantizing deep convolutional networks for efficient inference: A
615 whitepaper. *arXiv preprint arXiv:1806.08342*, 2018.
616
- 617 Changhun Lee, Jungyu Jin, Taesu Kim, Hyungjun Kim, and Eunhyeok Park. Owq: Outlier-aware
618 weight quantization for efficient fine-tuning and inference of large language models. In *Proceedings*
619 *of the AAAI Conference on Artificial Intelligence*, volume 38, pp. 13355–13364, 2024.
- 620 Min Li, Zihao Huang, Lin Chen, Junxing Ren, Miao Jiang, Fengfa Li, Jitao Fu, and Chenghua Gao.
621 Contemporary advances in neural network quantization: A survey. In *2024 International Joint*
622 *Conference on Neural Networks (IJCNN)*, pp. 1–10. IEEE, 2024a.
- 623 Shiyao Li, Xuefei Ning, Ke Hong, Tengxuan Liu, Luning Wang, Xiuhong Li, Kai Zhong, Guohao
624 Dai, Huazhong Yang, and Yu Wang. Llm-mq: Mixed-precision quantization for efficient llm
625 deployment. In *The Efficient Natural Language and Speech Processing Workshop with NeurIPS*,
626 volume 9, 2023.
627
- 628 Shiyao Li, Xuefei Ning, Luning Wang, Tengxuan Liu, Xiangsheng Shi, Shengen Yan, Guohao Dai,
629 Huazhong Yang, and Yu Wang. Evaluating quantized large language models. *arXiv preprint*
630 *arXiv:2402.18158*, 2024b.
- 631 Yixiao Li, Yifan Yu, Chen Liang, Nikos Karampatziakis, Pengcheng He, Weizhu Chen, and Tuo
632 Zhao. Loftq: Lora-fine-tuning-aware quantization for large language models. In *The Twelfth*
633 *International Conference on Learning Representations*, 2024c.
634
- 635 Ji Lin, Jiaming Tang, Haotian Tang, Shang Yang, Wei-Ming Chen, Wei-Chen Wang, Guangxuan
636 Xiao, Xingyu Dang, Chuang Gan, and Song Han. Awq: Activation-aware weight quantization for
637 on-device llm compression and acceleration. *Proceedings of Machine Learning and Systems*, 6:
638 87–100, 2024.
- 639 Zechun Liu, Barlas Oguz, Changsheng Zhao, Ernie Chang, Pierre Stock, Yashar Mehdad, Yangyang
640 Shi, Raghuraman Krishnamoorthi, and Vikas Chandra. Llm-qat: Data-free quantization aware
641 training for large language models. *arXiv preprint arXiv:2305.17888*, 2023.
642
- 643 Zirui Liu, Jiayi Yuan, Hongye Jin, Shaochen Zhong, Zhaozhuo Xu, Vladimir Braverman, Beidi
644 Chen, and Xia Hu. Kivi: A tuning-free asymmetric 2bit quantization for kv cache. *arXiv preprint*
645 *arXiv:2402.02750*, 2024.
- 646 Yuexiao Ma, Taisong Jin, Xiawu Zheng, Yan Wang, Huixia Li, Yongjian Wu, Guannan Jiang, Wei
647 Zhang, and Rongrong Ji. Ompq: Orthogonal mixed precision quantization. In *Proceedings of the*
AAAI conference on artificial intelligence, volume 37, pp. 9029–9037, 2023.

- 648 Stephen Merity, Caiming Xiong, James Bradbury, and Richard Socher. Pointer sentinel mixture
649 models. *arXiv preprint arXiv:1609.07843*, 2016.
- 650
- 651 Maxim Naumov, Utku Diril, Jongsoo Park, Benjamin Ray, Jędrzej Jablonski, and Andrew Tul-
652 loch. On periodic functions as regularizers for quantization of neural networks. *arXiv preprint*
653 *arXiv:1811.09862*, 2018.
- 654 Gunho Park, Minsub Kim, Sungjae Lee, Jeonghoon Kim, Beomseok Kwon, Se Jung Kwon, Byeong-
655 wook Kim, Youngjoo Lee, Dongsoo Lee, et al. Lut-gemm: Quantized matrix multiplication
656 based on luts for efficient inference in large-scale generative language models. In *The Twelfth*
657 *International Conference on Learning Representations*, 2024.
- 658
- 659 Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor
660 Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style,
661 high-performance deep learning library. *Advances in neural information processing systems*, 32,
662 2019.
- 663 Zhongnan Qu, Zimu Zhou, Yun Cheng, and Lothar Thiele. Adaptive loss-aware quantization for
664 multi-bit networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern*
665 *Recognition*, pp. 7988–7997, 2020.
- 666
- 667 Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi
668 Zhou, Wei Li, and Peter J Liu. Exploring the limits of transfer learning with a unified text-to-text
669 transformer. *Journal of machine learning research*, 21(140):1–67, 2020.
- 670 Mariam Rakka, Mohammed E Fouda, Pramod Khargonekar, and Fadi Kurdahi. Mixed-precision
671 neural networks: A survey. *arXiv preprint arXiv:2208.06064*, 2022.
- 672
- 673 Wenqi Shao, Mengzhao Chen, Zhaoyang Zhang, Peng Xu, Lirui Zhao, Zhiqian Li, Kaipeng Zhang,
674 Peng Gao, Yu Qiao, and Ping Luo. Omniquant: Omnidirectionally calibrated quantization for large
675 language models. *arXiv preprint arXiv:2308.13137*, 2023.
- 676 Romal Thoppilan, Daniel De Freitas, Jamie Hall, Noam Shazeer, Apoorv Kulshreshtha, Heng-Tze
677 Cheng, Alicia Jin, Taylor Bos, Leslie Baker, Yu Du, et al. Lamda: Language models for dialog
678 applications. *arXiv preprint arXiv:2201.08239*, 2022.
- 679
- 680 Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée
681 Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. Llama: Open and
682 efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023a.
- 683
- 684 Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay
685 Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. Llama 2: Open foundation
686 and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*, 2023b.
- 687
- 688 Albert Tseng, Jerry Chee, Qingyao Sun, Volodymyr Kuleshov, and Christopher De Sa. Quip#:
689 Even better llm quantization with hadamard incoherence and lattice codebooks. In *Forty-first*
690 *International Conference on Machine Learning*, 2024.
- 691
- 692 Kuan Wang, Zhijian Liu, Yujun Lin, Ji Lin, and Song Han. Haq: Hardware-aware automated
693 quantization with mixed precision. In *Proceedings of the IEEE/CVF conference on computer*
694 *vision and pattern recognition*, pp. 8612–8620, 2019.
- 695
- 696 Xiuying Wei, Yunchen Zhang, Xiangguo Zhang, Ruihao Gong, Shanghang Zhang, Qi Zhang, Fengwei
697 Yu, and Xianglong Liu. Outlier suppression: Pushing the limit of low-bit transformer language
698 models. *Advances in Neural Information Processing Systems*, 35:17402–17414, 2022.
- 699
- 700 Bichen Wu, Yanghan Wang, Peizhao Zhang, Yuandong Tian, Peter Vajda, and Kurt Keutzer. Mixed
701 precision quantization of convnets via differentiable neural architecture search. *arXiv preprint*
arXiv:1812.00090, 2018.
- 702
- 703 Mengzhou Xia, Tianyu Gao, Zhiyuan Zeng, and Danqi Chen. Sheared llama: Accelerating language
704 model pre-training via structured pruning. In *The Twelfth International Conference on Learning*
705 *Representations*, 2024.

- 702 Guangxuan Xiao, Ji Lin, Mickael Seznec, Hao Wu, Julien Demouth, and Song Han. Smoothquant:
703 Accurate and efficient post-training quantization for large language models. In *International*
704 *Conference on Machine Learning*, pp. 38087–38099. PMLR, 2023.
- 705
706 Guangxuan Xiao, Yuandong Tian, Beidi Chen, Song Han, and Mike Lewis. Efficient streaming
707 language models with attention sinks. In *The Twelfth International Conference on Learning*
708 *Representations*, 2024.
- 709 Yuhui Xu, Lingxi Xie, Xiaotao Gu, Xin Chen, Heng Chang, Hengheng Zhang, Zhengsu Chen,
710 XIAOPENG ZHANG, and Qi Tian. Qa-lora: Quantization-aware low-rank adaptation of large
711 language models. In *The Twelfth International Conference on Learning Representations*, 2024.
- 712
713 Zhewei Yao, Zhen Dong, Zhangcheng Zheng, Amir Gholami, Jiali Yu, Eric Tan, Leyuan Wang,
714 Qijing Huang, Yida Wang, Michael Mahoney, et al. Hawq-v3: Dyadic neural network quantization.
715 In *International Conference on Machine Learning*, pp. 11875–11886. PMLR, 2021.
- 716
717 Zhewei Yao, Reza Yazdani Aminabadi, Minjia Zhang, Xiaoxia Wu, Conglong Li, and Yuxiong
718 He. Zeroquant: Efficient and affordable post-training quantization for large-scale transformers.
719 *Advances in Neural Information Processing Systems*, 35:27168–27183, 2022.
- 720
721 Zhihang Yuan, Lin Niu, Jiawei Liu, Wenyu Liu, Xinggang Wang, Yuzhang Shang, Guangyu Sun,
722 Qiang Wu, Jiayang Wu, and Bingzhe Wu. Rptq: Reorder-based post-training quantization for
723 large language models. *arXiv preprint arXiv:2304.01089*, 2023.
- 724
725 Susan Zhang, Stephen Roller, Naman Goyal, Mikel Artetxe, Moya Chen, Shuohui Chen, Christopher
726 Dewan, Mona Diab, Xian Li, Xi Victoria Lin, et al. Opt: Open pre-trained transformer language
727 models. *arXiv preprint arXiv:2205.01068*, 2022.
- 728
729 Zhaoyang Zhang, Wenqi Shao, Jinwei Gu, Xiaogang Wang, and Ping Luo. Differentiable dynamic
730 quantization with mixed precision and adaptive resolution. In *International Conference on Machine*
731 *Learning*, pp. 12546–12556. PMLR, 2021.
- 732
733 Zixuan Zhou, Xuefei Ning, Ke Hong, Tianyu Fu, Jiaming Xu, Shiyao Li, Yuming Lou, Luning Wang,
734 Zhihang Yuan, Xiuhong Li, et al. A survey on efficient inference for large language models. *arXiv*
735 *preprint arXiv:2404.14294*, 2024.
- 736
737
738
739
740
741
742
743
744
745
746
747
748
749
750
751
752
753
754
755

A APPENDIX

A.1 IMPACT OF SPARSITY LEVELS OF CMPQ

As discussed in Section 4.6, both activation-based and quantization-based outliers contribute to performance improvements. To evaluate the trade-off between performance and outlier protection ratio, we adjusted the ratio of activation-based outliers from 0.05% to 0.45% and present the perplexity results of the 3-bit quantized OPT-6.7B model on the C4 benchmarks, with varying outlier extraction percentages ranging from 0% to 0.5%, as shown in Fig 4. Notably, while we maintain a fixed protection ratio of 0.5% for quantization-based outliers across all experiments to ensure fair comparisons, the plot reveals that the perplexity gains diminish when the protection ratio exceeds 0.2%. This finding highlights CMPQ’s potential to achieve superior performance with reduced storage requirements.

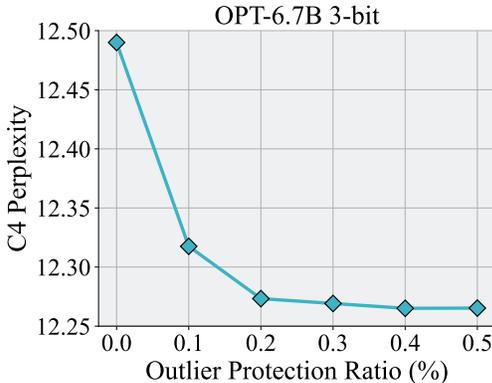


Figure 4: Outlier protection ratio and perplexity trade-off of 3-bit quantized OPT-6.7B model.

A.2 IMPACT OF NON-UNIFORM QUANTIZATION

In Table 7, we provide a detailed analysis to further clarify the impact of non-uniform quantization. For uniform quantization, we apply the widely used round-to-nearest method with a group size of 128 for channel-wise weight quantization, while preserving 0.5% of activation-based outliers to ensure a fair comparison. Additionally, we report the best perplexity achieved by the baseline methods. As shown in Table 7, across various bit-widths and model sizes, non-uniform quantization consistently outperforms uniform quantization, particularly in extremely low-bit (2-bit) settings. This is because the non-uniform distribution of weights leads to inefficient utilization of the quantization bins in uniform quantization, where some bins may remain underutilized or unused. Interestingly, we also observe that for certain tasks, uniform quantization can improve perplexity (e.g., OPT-13B at 3-bit on WikiText2). In such cases, equipping CMPQ with uniform quantization yields the best performance.

Table 7: Perplexity (\downarrow) comparison on the C4 and WikiText-2 datasets. LLMs are quantized by CMPQ using non-uniform and uniform approaches.

Method	Avg. Bit	OPT-2.7B		OPT-6.7B		OPT-13B	
		Wiki	C4	Wiki	C4	Wiki	C4
Baseline	2	>100	38.07	22.33	21.62	16.02	16.60
Uniform	2	>100	80.52	>100	>100	>100	>100
Non-Uniform	2	32.46	28.32	18.63	18.31	16.48	16.30
Baseline	3	16.32	15.63	11.41	12.30	10.50	12.39
Uniform	3	13.49	14.06	11.41	12.28	10.42	11.65
Non-Uniform	3	13.38	14.05	11.45	12.26	10.67	11.64
Baseline	4	12.69	13.48	10.93	11.86	10.21	11.28
Uniform	4	12.63	13.34	10.91	11.84	10.19	11.27
Non-Uniform	4	12.63	13.33	10.95	11.83	10.17	11.27

A.3 ADDITIONAL EXPERIMENTAL RESULTS

In Table 8, we present a comparison of quantization results across additional LLMs, including models from the OPT family ranging from 1.3B to 30B parameters, as well as the more recent LLaMA3-7B. Our findings are consistent with the main results, indicating that CMPQ generally outperforms all

810 baselines. However, we observe that as model size increases and the models become more powerful,
 811 the performance gap between different methods narrows. For instance, on the OPT-30B model,
 812 CMPQ outperforms baselines in only 2 out of 6 tasks. We attribute this to the differing quantization
 813 strategies: while the baselines quantize weights in blocks of 128, CMPQ performs quantization
 814 channel-wise. As model size grows, the additional bits allocated by these block-based methods may
 815 exceed those introduced by CMPQ, leading to similar performance outcomes. Nevertheless, CMPQ
 816 remains competitive in terms of performance while offering lower storage requirements.

817 Table 8: Additional perplexity (\downarrow) comparison of {2, 3, 4}-bit quantized LLMs on the C4 and WikiText-2
 818 datasets. Bold font indicates the best performance across all methods, while underlined results denote the
 819 second-best.
 820

Method	Avg. Bit	OPT-1.3B		OPT-13B		OPT-30B		LLaMA3-8B	
		Wiki	C4	Wiki	C4	Wiki	C4	Wiki	C4
FP16	16	14.62	14.72	10.13	11.2	9.56	10.69	6.1	9.2
RTN	2	>1000	>1000	>1000	>1000	>1000	>1000	>1000	>1000
GPTQ	2	>1000	>1000	372.68	135.48	71.7	29.59	210	>1000
AWQ	2	-	-	-	-	-	-	>1000	>1000
QuIP	2	<u>41.64</u>	29.78	16.02	<u>16.6</u>	11.48	<u>13.55</u>	85.1	<u>130</u>
CMPQ	2	41.58	<u>36.67</u>	<u>16.48</u>	16.3	<u>11.53</u>	12.89	<u>120</u>	110
RTN	3	>1000	>1000	>1000	>1000	>1000	>1000	27.9	110
GPTQ	3	21.35	21.59	11.6	13.34	10.32	12.23	8.2	13.7
AWQ	3	16.32	<u>16.28</u>	10.67	12.61	9.85	10.96	8.2	11.6
QuIP	3	<u>16.21</u>	<u>17.12</u>	10.5	<u>12.39</u>	9.79	11.66	7.5	<u>11.3</u>
CMPQ	3	16.07	16.08	<u>10.67</u>	11.64	<u>9.83</u>	<u>10.98</u>	<u>7.89</u>	11.3
RTN	4	47.62	27.2	11.32	12.35	10.77	13.52	8.5	13.4
GPTQ	4	15.59	16.96	10.31	12.26	9.63	11.8	7.0	10.4
AWQ	4	14.94	<u>15.04</u>	10.22	<u>11.28</u>	9.59	<u>10.75</u>	6.6	<u>9.4</u>
QuIP	4	<u>14.88</u>	<u>16.38</u>	<u>10.21</u>	<u>12.16</u>	9.61	11.5	6.6	11.3
CMPQ	4	14.84	14.99	10.17	11.27	<u>9.61</u>	10.74	6.5	9.39