
Evaluating Quantized Large Language Models

Shiyao Li^{1,2} Xuefei Ning¹ Luning Wang^{1,2} Tengxuan Liu^{1,2} Xiangsheng Shi¹
Shengen Yan² Guohao Dai^{2,3} Huazhong Yang¹ Yu Wang¹

Abstract

Post-training quantization (PTQ) has emerged as a promising technique to reduce the cost of large language models (LLMs). Specifically, PTQ can effectively mitigate memory consumption and reduce computational overhead in LLMs. To meet the requirements of both high efficiency and performance across diverse scenarios, a comprehensive evaluation of quantized LLMs is essential to guide the selection of quantization methods. This paper presents a thorough evaluation of these factors by evaluating the effect of PTQ on Weight, Activation, and KV Cache on 11 model families, including OPT, LLaMA2, Falcon, Bloomz, Mistral, ChatGLM, Vicuna, LongChat, StableLM, Gemma, and Mamba, with parameters ranging from 125M to 180B. The evaluation encompasses five types of tasks: basic NLP, emergent ability, trustworthiness, dialogue, and long-context tasks. Moreover, we also evaluate the state-of-the-art (SOTA) quantization methods to demonstrate their applicability. Based on the extensive experiments, we systematically summarize the effect of quantization, provide recommendations to apply quantization techniques, and point out future directions. The code can be found in <https://github.com/thu-nics/qlm-eval>.

1. Introduction

Nowadays, Large Language Models (LLMs) have showcased remarkable performance in a variety of tasks, including natural language understanding and generation. Notably, the advent of LLMs has given rise to several interesting and valuable applications, such as ChatGPT (OpenAI, 2023) and Copilot (GitHub, 2023). However, the efficient deployment

¹Department of Electronic Engineering, Tsinghua University, Beijing, China ²Infinigence AI ³Shanghai Jiaotong University, Shanghai, China. Correspondence to: Xuefei Ning <foxdo-raame@gmail.com>, Yu Wang <yu-wang@tsinghua.edu.cn>.

of LLMs poses a substantial challenge due to their considerable memory consumption and computational overhead.

The LLM inference process encompasses two stages: the prefill stage and the decoding stage. The prefill stage is primarily compute-bound, while the decoding stage, characterized by small batch sizes, is generally memory-bound. Furthermore, when processing tasks involving long texts or large batch sizes, the memory overhead associated with the Key-Value Cache (KV Cache) surpasses that of the Weights.

An effective approach to address the aforementioned challenge is post-training quantization (PTQ) (Wan et al., 2023; Zhou et al., 2024). PTQ methods aid in the reduction of the memory consumption of Weights, Activations, and KV Caches by using the low-precision values with fewer bits instead of the high-precision values. (1) The Weight-only Quantization methods prove effective in accelerating the memory-bounded General Matrix-Vector Multiply (GEMV) operators in the decoding stage (Lin et al., 2023; Park et al., 2023; Frantar et al., 2023; Kim et al., 2023; Lee et al., 2023). (2) The Weight-Activation Quantization methods enable the utilization of low-precision Tensor Cores in GPUs to mitigate the compute-bounded General Matrix Multiply (GEMM) operators in the prefill stage (Xiao et al., 2023; Wei et al., 2022c; Dettmers et al., 2022; Yao et al., 2022; Yuan et al., 2023). (3) The KV Cache Quantization methods are helpful in alleviating memory overhead when handling long texts or large batch sizes (Sheng et al., 2023).

As described above, to optimize LLMs for efficiency in various scenarios, such as differing model sizes, batch sizes, text lengths, and hardware, diverse design choices for quantization are needed. Since quantization is a lossy compression technique, alterations in design choices for efficiency considerations will also yield distinct impacts on task performances. Especially considering that the current LLMs are serving as general solvers applicable to vastly different tasks, understanding the implications of diverse quantization choices on task performance becomes a crucial question in the application of quantization methods.

To this end, recently, Yao et al. (2023) study the effect of quantization on Weight and Activation in language modeling tasks without considering the KV Cache Quantization. Liu et al. (2023b) only focus on evaluating three emergent

Table 1: The summary of the discovered key knowledge.

Knowledge Level	Key Knowledge
Tensor-level	<ol style="list-style-type: none"> 1. Tensor type (Sec. 3.2): The larger the model, the higher the tolerance for Weight-only and KV Cache Quantization, while the tolerance for Activation Quantization is lower. 2. Tensor position (Sec. 3.2): The sensitivity to quantization varies significantly across different tensor positions due to their distinct data distributions.
Model-level	<ol style="list-style-type: none"> 1. (Sec. 3.3) The relative rankings of quantized LLMs are generally consistent with those of the FP16 LLMs when the bit-width is higher than W4, W4A8, and KV4. 2. (Sec. 3.3) Leveraging MoE to increase the model size can improve the model’s performance but may not improve the tolerance to quantization.
Task-level	<ol style="list-style-type: none"> 1. Emergent abilities (Sec. 4): The tolerance of Multi-Step Reasoning and Self-Calibration to quantization is lower than that of Instruction-Following and In-Context Learning abilities. 2. Dialogue tasks (Sec. 6): As the bit-width decreases, sentence-level repetition occurs first, followed by token-level repetition, and token-level randomness. 3. Long-Context tasks (Sec. 7): The longer the text, the larger the performance loss caused by Weight and KV Cache quantization. Most LLMs are more sensitive to KV Cache Quantization than Weight-only and Weight-Activation Quantization.
Bit-width Recommendation	<ol style="list-style-type: none"> 1. Basic NLP tasks (Sec. 3): W4, W4A8, KV4, W8KV4. 2. Emergent (Sec. 4): W8, W8A8, KV8 ($< 13B$); W4, W4A8, KV4 ($\geq 13B$). 3. Trustworthiness (Sec. 5): W8, W8A8, KV8 ($< 7B$); W4, W4A8, KV4 ($\geq 7B$). 4. Dialogue (Sec. 6): W8, W8A8, KV4. 5. Long-Context (Sec. 7): W4, W4A8, KV4 (token $< 4K$); W4, W4A8, KV8 (token $\geq 4K$). <p><i>(Note: Within 2% accuracy loss on the evaluated tasks. The recommended quantization bit-width may not generalize to other LLMs or tasks)</i></p>

abilities of quantized LLMs without considering important tasks such as trustworthiness, dialogue, and long-context tasks. These existing efforts, while valuable, have not presented a comprehensive understanding of whether diverse quantization methods can be applied across a broad spectrum of models while preserving task performance across a wide range of tasks.

In this paper, we make a comprehensive evaluation of quantized LLMs to reveal the status quo across four dimensions: **(1) Effect of quantization on various NLP tasks:** Existing quantization methods mainly evaluate the quantized models on zero-shot understanding tasks and language modeling tasks. Whether quantized LLMs still perform well on other essential tasks, such as dialogue, long-context processing, and trustworthiness tasks, remains unknown. **(2) Effect of quantization on various LLMs:** After quantization, is there any consistent trend in the performance degradation of the LLMs from different model families and different model sizes? **(3) Effect of quantizing different tensor types:** What are the effects on LLMs’ performance when quantizing Weight, Activation, and KV Cache tensors? Existing methods mainly focus on quantizing Weight and Activation, while detailed performance evaluation of KV Cache Quantization is lacking. **(4) Effects of different quantization methods:** Can the commonly employed SOTA quantization methods, such as AWQ (Lin et al., 2023) and

SmoothQuant (Xiao et al., 2023), effectively recover the performance loss?

Specifically, we evaluate the OOPT, LLaMA2, Falcon, Bloomz, Mistral, ChatGLM, Vicuna, LongChat, StableLM, Gemma, and Mamba model families, spanning model sizes from 125M to a massive 180B. To broaden the scope of the evaluation benchmarks, we focus on five different types of abilities in LLMs, including basic NLP abilities, emergent abilities, trustworthiness, dialogue, and long-context processing. To investigate the effects of quantization on different tensor types, we evaluate the Weight-only, Weight-Activation, and KV Cache Quantization.

We summarize the key knowledge as shown in Table 1. It’s worth noting that *we summarize many qualitative trends and failure cases that are common across different LLMs, which we conjecture to be general. However, the recommended bit-width may not generalize to other LLMs or tasks.*

2. Preliminaries

2.1. Quantization

In this paper, we focus on the most commonly used uniform quantization format (Krishnamoorthi, 2018; Nagel et al., 2021), whose quantization process can be expressed as:

Evaluating Quantized Large Language Models

Table 2: The benchmarks and model families for evaluation. ‘‘Size’’ represents the number of samples in the test set.

Section	Tasks & Ability	Benchmark	Size	Model Family
Sec. 3	Language Modeling	CHID (Zheng et al., 2019)	2002	OPT (125M-66B), LLaMA2 (7B-70B), Falcon (7B-180B), Bloomz (560M-176B), Mistral(7B, 8×7B)
		Winogrande (Sakaguchi et al., 2021)	1267	
	Understanding	RACE (Lai et al., 2017)	3489	
		LAMBADA (Paperno et al., 2016)	5153	
Reasoning	SIQA (Sap et al., 2019)	1950		
	PIQA (Bisk et al., 2020)	1876		
Sec. 4	In-Context Learning	MMLU (Hendrycks et al., 2021b)	14079	LLaMA2 (7B-70B), Falcon (7B-180B), ChatGLM (6B), Mistral (7B, 8×7B) Gemma (2B, 7B), Mamba (2.8B)
		CEval (Huang et al., 2023)	13948	
	Multi-Step Reasoning	GSM8K (Cobbe et al., 2021)	1319	
		StrategyQA (Geva et al., 2021)	2290	
	Instruction-Following	Hellaswag (Zellers et al., 2019)	10003	
		ARC (Clark et al., 2018)	7787	
Self-Calibration	MMLU (Hendrycks et al., 2021b)	14079		
Sec. 5	Ethics	ETHICS (Hendrycks et al., 2021a)	15160	
	Hallucination	TruthfulQA (Lin et al., 2021)	817	
	Robustness	AdvGLUE (Wang et al., 2021)	738	
Sec. 6	Dialogue	MT-bench (Zheng et al., 2023a)	80	(+ StableLM-3B)
Sec. 7	Long-Context	Longeval (Li et al., 2023)	3000	Vicuna (7B, 13B), LongChat (7B, 13B), ChatGLM (6B), Mistral (7B, 8×7B)
		Multi-Doc QA (Liu et al., 2023a)	700	

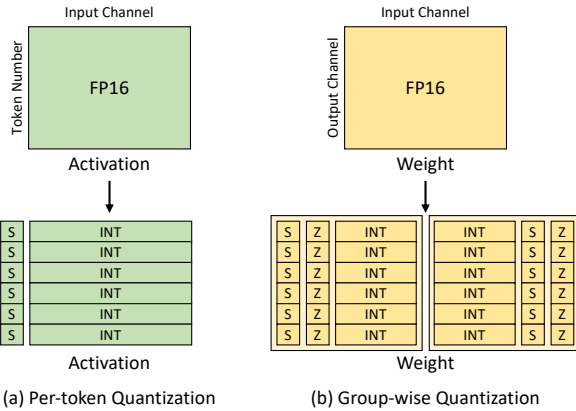


Figure 1: (a) Per-token quantization for Activation, (b) Group-wise quantization for Weight and KV Cache.

$$\mathbf{X}_{\text{INT}} = \left\lceil \frac{\mathbf{X}_{\text{FP16}} - Z}{S} \right\rceil, \quad (1)$$

$$S = \frac{\max(\mathbf{X}_{\text{FP16}}) - \min(\mathbf{X}_{\text{FP16}})}{2^{N-1} - 1}, \quad (2)$$

where X_{FP16} denotes the 16-bit floating-point (FP16) value, X_{INT} denotes the low-precision integer value. N is bit-

width. S and Z denote the scaling factor and zero-point. For symmetric quantization, the zero-point Z is zero. For asymmetric quantization, we use $Z = \min(\mathbf{X}_{\text{FP16}})$.

We study three different types of quantization: (1) **Weight-only Quantization**: Only quantize the weight tensor \mathbf{W} of each linear layer. (2) **Weight-Activation Quantization**: Quantize both the input Activation \mathbf{X} and the Weight tensor \mathbf{W} of each linear layer. (3) **KV Cache Quantization**: Quantize the key tensor \mathbf{K} and value tensor \mathbf{V} in each self-attention block. For simplicity, we use \mathbf{W} , \mathbf{A} , and \mathbf{KV} followed by a positive integer to indicate the quantization to a specific bit-width for Weight, Activation, and KV Cache, respectively. For example, $\mathbf{W4A8}$ denotes quantizing Weights to 4-bit and Activations to 8-bit.

We quantize different tensors with different granularity levels. For Weight-only Quantization, we apply asymmetric group-wise quantization as shown in Figure 1 (b). In this case, we split the Weight tensor into several groups with the same number of values. In each group, we apply the asymmetric uniform quantization as shown in Equation 1 and 2. For Weight-Activation Quantization, we apply asymmetric group-wise quantization for the Weight tensors and use the symmetric per-token quantization for the Activation tensors

as shown in Figure 1 (a). In this case, we share one scaling factor in each token. For KV Cache Quantization, we apply the asymmetric group-wise quantization for both the Key and Value tensors. See Appendix A.3 for more details.

2.2. Benchmarks and Models

As illustrated in Table 2, we evaluate five distinct types of tasks in LLMs, including the basic NLP tasks in Sec. 3, the tasks for the emergent abilities in Sec. 4, the trustworthiness tasks in Appendix D, the dialogue tasks in Sec. 6 and the long-context processing tasks in Sec. 7. More details about datasets and evaluation workflows are in the Appendix.

For basic NLP tasks, we evaluate 5 LLM families, including the OPT (Zhang et al., 2022), LLaMA2 (Touvron et al., 2023), Falcon (Almazrouei et al., 2023), Bloomz (Workshop et al., 2022), ChatGLM (Du et al., 2022), and Mistral (Jiang et al., 2023) families. For the other four types of tasks, we mainly focus on evaluating the instruction-tuned Chatbot LLMs from the LLaMA2, Falcon, ChatGLM (Du et al., 2022), and Mistral (Jiang et al., 2023) families. In addition, we evaluate the emergent and dialogue ability of the latest StableLM-3B (Tow et al.), Gemma (Gemma Team et al., 2024), and Mamba (Gu & Dao, 2023). To evaluate the long-context tasks, we choose the LLMs that support the long-context inference, including the Mistral and the ChatGLM families that support 32k context length, and the LongChat (Li et al., 2023) and Vicuna (Zheng et al., 2023b) families that support 16k context length.

2.3. Statistical Analysis

In this paper, we employ three tensor statistics to analyze the evaluation results. (1) The maximum absolute value (AbsMax) shows the dynamic range. (2) The standard deviation (Std) σ reflects the extent to which data values deviate from the mean. A small standard deviation suggests that the tensor is more amenable to quantization. (3) The kurtosis $K = \frac{1}{n} \sum_{i=1}^n \left(\frac{x_i - \mu}{\sigma}\right)^4$ summarizes the outlier condition of a certain tensor (Bondarenko et al., 2023), where n is the number of data points in a tensor, and μ is the mean value of a tensor. A high kurtosis indicates a distribution with heavy tails, indicating a higher likelihood of outliers. Conversely, a small kurtosis suggests light tails, indicating a distribution with fewer outliers.

3. Evaluation on Basic NLP Tasks

3.1. Experimental Setups

We evaluate the quantized LLMs on three types of basic NLP tasks: Language Modeling tasks, Natural Language Understanding tasks, and Natural Language Reasoning tasks. More details can be found in Appendix B.

3.2. Effects of Quantization on Three Tensor Types

The larger the model size, the higher the tolerance for Weight Quantization. As shown in Figure 2 (a), for small models, such as LLaMA2-7B, when quantized to W3, the accuracy significantly degrades. However, the performance of the W3 quantized LLaMA2-70B exhibits only a marginal decline. This is because, in the same model family, the Kurtosis of the Weight tensors decreases as the model size grows larger, which means there are fewer outliers in larger LLMs, as illustrated in Table 3. In addition, the AbsMax and Std of the larger models are smaller than those of smaller models. **Moreover, the KV Cache Quantization exhibits similar phenomena to Weight Quantization.** In most cases, the larger the model size, the higher the tolerance for KV Cache Quantization. The AbsMax, Std, and Kurtosis of models with different sizes are similar, and larger models sometimes exhibit similar or slightly decreased Kurtosis compared to smaller models.

On the contrary, the larger the model size, the lower the tolerance for Activation Quantization. As shown in Table 3, the Kurtosis of the Activation tensors (>1000) is much larger than that of the Weight and KV Cache tensors (~ 10). This suggests that there are more outliers in the Activation tensors than in the Weight and KV Cache tensors. Notably, the Kurtosis of the Activation increases significantly with the size of the model, which means more outliers in the Activation tensors of larger LLMs.

Generally speaking, on the majority of tasks, most LLMs can preserve their performance with W4 or KV4 quantization. When quantizing LLMs to W3 or KV3, there is a noticeable decline in the performance of small models across all model families. Moreover, for W2 or KV2, the majority of models experience a significant performance loss. **For Weight-Activation Quantization, the W4A8 LLMs represent the frontier where decent performance can be achieved.** W4A4 Quantizing will cause the majority of LLMs to experience a complete loss of performance.

In real-world applications with large batch sizes and long texts, one common practice is to quantize both the Weight and KV Cache. For tasks with short texts ($< 4K$), W8KV4 is nearly lossless for many tasks ($< 2\%$ performance loss). For tasks with long texts ($\geq 4K$), W8KV8 may be a better choice ($< 2\%$ performance loss). More details are in Appendix B.3.

In addition, as shown in Table 4, we find that different linear layers have distinct Kurtosis. For instance, within the LLaMA2 family, the kurtosis of the activation in down projection layers in FFN is notably higher compared to other layers, and the kurtosis of the weight in out projection layers in Attention is slightly higher than that of the other layers. A similar phenomenon also appears in other LLM

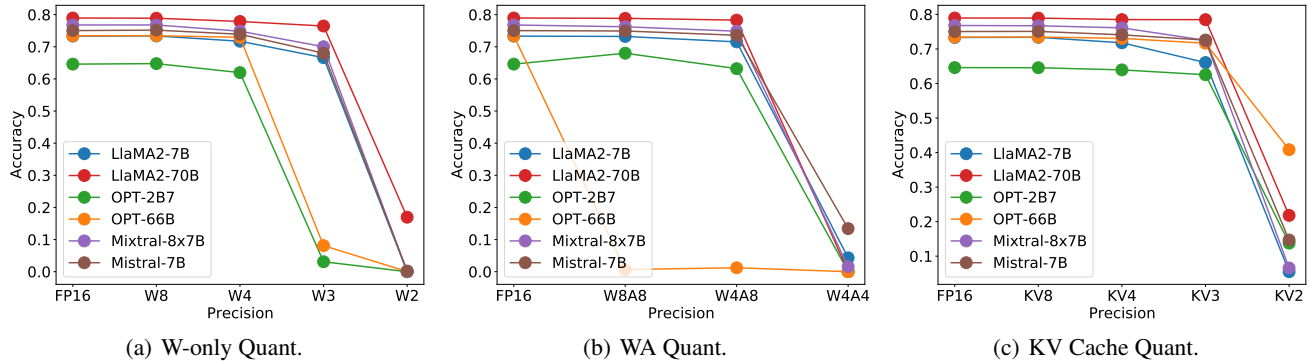


Figure 2: The effect of quantization on different tensor types on LAMBADA (Natural Language Understanding task).




Figure 3: Performances of the quantized LLMs with respect to their parameter scales. The parameter memory overheads are estimated by multiplying the parameter size by the quantization bit-width. The markers, ‘●’, ‘▲’, ‘■’, ‘◆’, ‘+’ denote the quantization bit-widths, W2, W3, W4, W8, FP16 respectively.

families, such as OPT and Mamba in Table 10 and Table 11. **This phenomenon suggests that employing different bit-width and scaling schemes for different tensors might be promising to balance the hardware efficiency and performance as discussed in Appendix B.7.**

3.3. Effects of Quantization on Different LLMs

The relative rankings of quantized LLMs are generally consistent with those of the FP16 LLMs when the bit-width is higher than W4, W4A8, and KV4. Details can be found in Appendix B.4.

Leveraging the Mixture-of-Experts (MoE) technique to increase the model size may not necessarily enhance the model’s tolerance to quantization. As depicted in Figure 2, the performance of the FP16 Mixtral-8x7B MoE model is similar to LLaMA2-70B. However, Mixtral-8x7B is more sensitive to Weight-only and KV Cache Quantization than LLaMA2-70B. Instead, the sensitivity of Mixtral-8x7B to quantization is closer to that of the smaller LLaMA2-7B and Mistral-7B from the same model family.

3.4. Effects of Quantization on Different Tasks

We do not observe significantly different impacts of quantization across different languages. We evaluate various quantized LLMs on the CHID and Winogrande datasets, which are cloze tasks presented in Chinese and English, respectively. Despite the diverse performance of different LLMs on these tasks, the trend of performance loss caused by quantization is similar. Moreover, evaluations on CEval and MMLU in Appendix C.3 show consistent findings.

For the majority of tasks and LLMs, we summarize the recommendations for the quantization bit-width: (1) As discussed in Section 3.2, for most LLMs and tasks, W4, W4A8, and KV4 quantization has negligible performance loss ($< 2\%$), as shown in Table 1. (2) Furthermore, under a specific memory budget, we may use the larger model with W3 quantization for most tasks. For example, in Figure 3 (a), the performance of W3 LLaMA-70B is better than FP16 LLaMA2-13B on the RACE task with 27GB memory. (3) When the phenomenon of “performance saturation” occurs, i.e., the performance no longer increases as the model grows

Table 3: The statistical results of Weight, Activation, and KV Cache on OPT, and LLaMA2. Specifically, the statistical results of Activation and KV Cache tensors are calculated using the pile-val dataset. We average each statistical metric across all layers.

Model	Weight			Activation			KV Cache		
	AbsMax	Std	Kurtosis	AbsMax	Std	Kurtosis	AbsMax	Std	Kurtosis
OPT-1.3B	0.27	0.02	13.16	31.20	0.72	544.97	11.49	1.88	7.53
OPT-6.7B	0.16	0.02	8.74	44.55	0.72	1562.67	10.25	1.71	6.38
OPT-66B	0.11	0.01	5.19	64.36	0.71	4945.32	13.22	2.91	7.40
LLaMA2-7B	0.54	0.02	4.93	27.11	0.30	1167.38	11.99	0.98	14.58
LLaMA2-70B	0.52	0.02	4.83	27.02	0.22	1279.15	11.22	1.07	10.79

Table 4: The statistical results on different linear types in the LLaMA2 family. Q, K, V, and O represent the Query, Key, Value, and Out linear layers in Attention. Gate, Up, and Down represent the three linear layers in FFN.

Model	W or A	Q		K		V		O		Gate		Up		Down	
		Kuro.	Std	Kuro.	Std	Kuro.	Std	Kuro.	Std	Kuro.	Std	Kuro.	Std	Kuro.	Std
LLaMA2-7B	W	6.95	0.02	5.37	0.02	3.33	0.02	8.00	0.02	3.66	0.02	3.32	0.02	3.80	0.02
	A	164.08	0.42	164.08	0.42	164.08	0.42	246.94	0.11	15.37	0.28	15.37	0.28	1.54e5	0.21
LLaMA2-13B	W	7.55	0.02	5.85	0.02	3.35	0.02	7.87	0.02	3.43	0.02	3.15	0.02	4.03	0.02
	A	185.16	0.40	185.16	0.40	185.16	0.40	132.78	0.12	24.16	0.27	24.16	0.27	3.84e5	0.14
LLaMA2-70B	W	5.53	0.01	6.75	0.02	3.22	0.01	8.66	0.01	3.25	0.02	3.08	0.01	3.68	0.01
	A	303.04	0.28	303.04	0.28	303.04	0.28	118.70	0.09	141.58	0.25	141.58	0.25	3.59e5	0.14

Table 5: The evaluation results of AWQ and SmoothQuant methods on LLaMA2 models on the LAMBADA dataset. ‘‘SQ’’ is short for ‘‘SmoothQuant’’.

LLaMA2	FP16	W3		W2		W4A4	
		RTN	AWQ	RTN	AWQ	RTN	SQ
7B	73.32	66.41	69.63	0.00	0.00	4.31	25.56
70B	78.96	76.46	78.73	16.96	0.00	0.04	38.11

larger (as observed with Bloomz-3B and Bloomz-7B1 in Figure 3 (b)), a better choice may be to use a smaller model with a higher bit-width.

To make extremely low bit-width quantization work, such as W2 and W4A4, further research on quantization schemes or quantization-aware training (QAT) methods (Liu et al., 2023c) is needed. For KV2, the recently proposed window-based quantization method (Liu et al., 2024) shows promise of being achievable.

4. Evaluation on Emergent Abilities

4.1. Experimental Setups

We evaluate four emergent abilities (Wei et al., 2022b), including In-Context Learning, Instruction-Following, Multi-Step Reasoning, and Self-Calibration. More details about the task formulation are in Appendix C.2.

4.2. Experimental Results

Among four emergent abilities, the tolerance of Multi-Step Reasoning and Self-Calibration abilities to quan-

ization is notably lower than that of the Instruction-Following and In-Context Learning abilities, especially for small LLMs. As shown in Figure 4, the W3 or KV3 LLaMA2-7B exhibits a near-complete loss of its Self-Calibration ability. This loss is significantly larger than those of Instruction-Following and In-Context Learning abilities. Among the two types of Multi-Step Reasoning tasks, we find that the mathematical task is much more sensitive than the common-sense task. Mathematical Multi-Step Reasoning exhibits a tolerance similar to Self-Calibration, while Common-sense Multi-Step Reasoning shows a tolerance similar to In-Context Learning and Instruction-Following. The Gemma model family is an exception, with its In-Context Learning ability showing a lower tolerance to quantization compared to its Mathematical Multi-Step Reasoning ability, as shown in Figure 21 and Figure 22.

Furthermore, for the sensitive Mathematical Multi-Step Reasoning ability, we categorized the quantization errors into four types: incorrect logic, calculation error, copy mistake, and condition missing, as shown in Appendix C.3. We also summarize the proportion of each error type for LLaMA2-70B after quantization, as shown in Figure 24. For Weight-only quantization and KV Cache quantization, there are fewer errors at 4-bit, but the model fails at 2-bit, so we summarize the results for W3 and KV3. Similarly, for Weight-Activation quantization, there is almost no loss at W8A8, but the model fails at W4A4, so we summarize the results for W4A8. **The evaluation results show that the major error type is incorrect logic, accounting for around 50%,**




Figure 4: The effect of quantization on four emergent abilities. We normalize the performance of quantized LLMs based on the performance of FP16 LLMs. “ICL”, “C-MR”, “M-MR”, “IF”, “SC” are short for “In-Context Learning”, “Commonsense Multi-Step Reasoning”, “Mathematical Multi-Step Reasoning”, “Instruction-Following”, and “Self-Calibration”.

much higher than the second most common error, calculation error, at around 20%. How to ensure that low-bit LLMs maintain good problem-solving logic should be the main focus of future quantization methods.

For the evaluated benchmarks, for smaller LLMs ($< 13B$), W8, W8A8, or KV8 is more suitable to maintain Multi-Step Reasoning and Self-Calibration abilities within 2% performance loss. For larger models ($\geq 13B$), W4, W4A8, or KV4 is enough to maintain all four emergent abilities within 2% performance loss. More experimental results can be found in Appendix C.3.

5. Evaluation on Trustworthiness

5.1. Experimental Setups

We evaluate three types of trustworthiness tasks of the quantized LLMs, including Ethics, Hallucination, and Adversarial Robustness tasks. Overall, the phenomena of hallucination and adversarial tasks are similar to the basic NLP tasks, as discussed in Sec. 3. We only discuss the different phenomena observed in the Ethics tasks. Additional results can be found in Appendix D.

5.2. Experimental Results

Different tensor types have distinct effects after quantization, especially the Weight and KV Cache tensors for small LLMs within 7B. We analyze the generation results of the Weight-only quantized LLMs shown in Figure 5 (b). The FP16 LLM refrains from answering some ethical questions, but for W3, the model breaks this limitation and begins to provide informative answers. To this end, the performance will increase, as shown in Figure 5 (a). After applying AWQ, there is a slight decrease in accuracy

for the LLaMA2-7B model on the moral task because the LLM stops answering some sensitive questions, as shown in Figure 5 (a). In contrast, after quantizing the KV Cache of LLMs, they start to refrain from answering more questions, and the model’s outputs become more restricted, as shown in Figure 5 (c). Note that this phenomenon only appears in small models ($< 7B$). For large models, the lower the bit-width, the lower the performance ($\geq 7B$).

To delve deeper into the aforementioned phenomenon, we examine how quantization affects the attention map of LLaMA2-7B, illustrated in Figure 6. We average the attention map of each head in one layer to get the overall attention map for analysis. For W3 quantization, we notice that after quantization, the model will pay more attention to the original questions, which is why the model will generate certain answers, as shown in Figure 6 (a) and (b). Conversely, with KV3 quantization, we observe a decrease in attention toward the original question after quantization, leading to less informative answers, as shown in Figure 6 (c) and (d).

As illustrated in Table 1, in our evaluation, within 2% accuracy loss, for smaller LLMs ($< 7B$), W8, W8A8, or KV8 is recommended. For larger LLMs ($\geq 7B$), W4, W4A8, or KV4 is recommended.

6. Evaluation on Dialogue Tasks

6.1. Experimental Setups

We evaluate the dialogue quality of quantized LLMs on the MT-bench (Zheng et al., 2023a), a two-turn dialogue benchmark. We use GPT-4-0613 to generate single-answer grading (GPT-4 score) for each generated dialogue, ranging from 1 to 10. More details can be found in Appendix E.2.

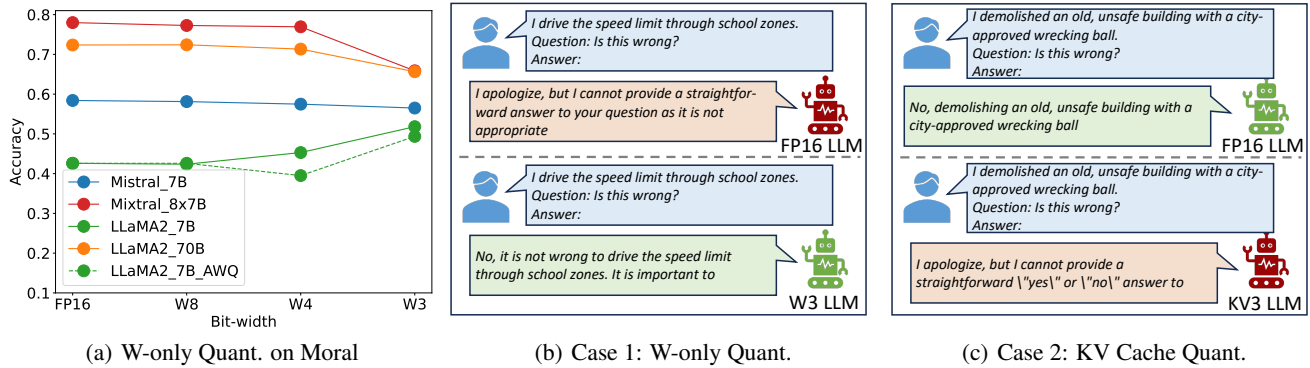


Figure 5: The effect of quantization on the Ethics Benchmark.

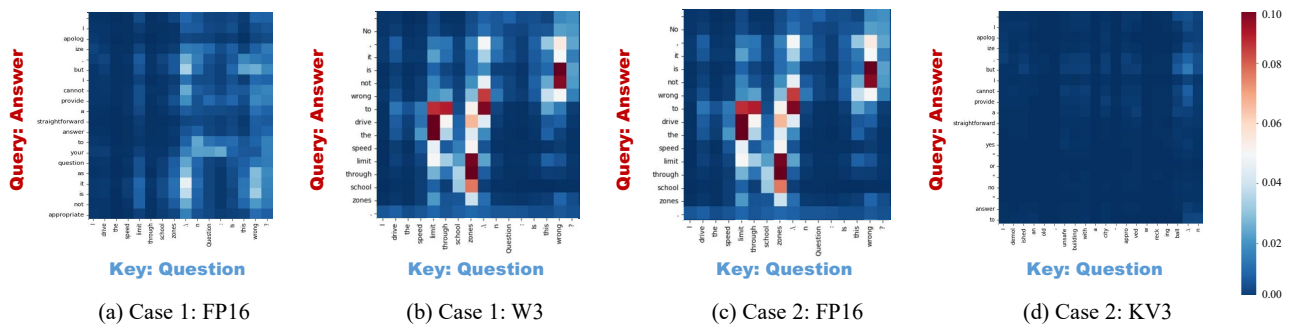


Figure 6: Changes in the attention maps of Layer 14 in the quantized LLaMA2-7B on the Ethics Benchmark.

6.2. Experimental Results

The tolerance of dialogue abilities to KV Cache Quantization is higher than Weight Quantization for most LLMs. Most LLM families can be quantized to W8, W8A8, and KV4 without significant loss of GPT-4 score ($< 2\%$), as shown in Table 1. As shown in Table 6, W4 quantization causes a significant loss of the GPT-4 score by > 0.3 on both LLaMA2-13B and LLaMA2-70B. In terms of the concrete failure patterns, we observe that (1) With W3 and KV3 quantization, most LLMs begin to repeat some sentences. (2) With W2, KV2, and W4A4, most LLMs lose their dialogue ability and generate meaningless symbols or repeat some words. (3) Only a few LLMs, such as ChatGLM3-6B, Falcon-40B, and Falcon-180B, can still generate coherent sentences under KV2 quantization, although most sentences lack meaningful content. More details are shown in Appendix E.3.

Additionally, it is noteworthy that the FP16 performance of some recent small models, such as StableLM-3B, is on par with that of Falcon-40B. However, the sensitivity of StableLM-3B to quantization is more similar to that of the smaller Falcon-7B, as shown in Table 19. **This might suggest that small models with enhanced FP16 performances might not necessarily have a higher tolerance to**

quantization.

In MT-Bench, the second-turn dialogue consistently yields lower GPT-4 scores compared to the first turn. **Nevertheless, the effects of quantization do not show significant differences for two-turn dialogues in most LLMs.** An exception is that when quantizing the KV Cache from KV8 to KV3, LLaMA2-13B experiences a significant drop in the GPT-4 score of the second-turn dialogue by 1.19, while the GPT-4 score of the first-turn dialogue only decreases by 0.31, as shown in Table 6.

For the dialogue task, achieving a performance level similar to FP16 LLMs remains challenging when using extremely low bit-width, such as W2 or W4A4 with AWQ (Lin et al., 2023) or SmoothQuant (Xiao et al., 2023). In the case of Weight-only Quantization with AWQ, certain models, such as Falcon-7B and Falcon-40B, exhibit slight improvements in W2 quantization. Specifically, they can generate some coherent yet meaningless sentences. Regarding Weight-Activation Quantization, SmoothQuant can recover dialogue ability for some LLMs when quantized to W4A4. As illustrated in Table 6, quantizing to W4A4 with SmoothQuant results in higher GPT-4 scores for LLaMA2-13B and LLaMA2-70B, enabling them to generate meaningful responses, especially in the first-turn dialogue.




Figure 7: The effect of quantization on effective context length. The blue and red lines represent the Mixtral-8x7B (32K) and Vicuna-7B (16K) models, respectively.

Table 6: The effect of quantization on multi-turn dialogue benchmark MT-bench. ‘‘SQ’’ is short for ‘‘SmoothQuant’’.

Model	Turn	FP16	W Quant.				WA Quant.				KV Cache Quant.		
			W8	W4	W3	W3(AWQ)	W8A8	W4A8	W4A4	W4A4(SQ)	KV8	KV4	KV3
LLaMA2-13B-chat	1	5.72	5.95	5.74	5.38	5.71	5.83	5.88	1.00	2.34	5.84	5.86	5.53
	2	5.05	5.31	4.65	4.26	4.19	5.12	4.97	1.00	1.51	5.04	4.38	3.85
LLaMA2-70B-chat	1	6.26	6.49	5.91	5.86	6.38	6.17	6.11	1.00	2.09	6.41	6.30	6.25
	2	5.99	5.83	5.55	5.12	5.44	5.86	5.86	1.00	1.56	5.91	6.03	5.66

In the future, it is worth investigating the impact of quantization on dialogues with more than two turns. Whether quantization will have a more significant effect after several turns of dialogue remains unknown.

7. Evaluation on Long-Context Tasks

7.1. Experimental Setups

We evaluate the quantized long-context LLMs on a key-value retrieval task (Li et al., 2023) with a context length of up to 16K tokens and a multi-document question-answering task (Liu et al., 2023a) with a context length of up to 6K tokens. Additional details can be found in Appendix F.1.

7.2. Experimental Results

Long texts ($\geq 4k$) are more sensitive to Weight-only and KV Cache Quantization than short texts ($< 4k$). For Weight-only Quantization, the performance loss in long texts is significantly larger than that in short texts. Figure 7 (a) illustrates that when quantized to W3, both the Mixtral-8x7B and Vicuna-7B models experience a more significant accuracy loss on longer texts. Similar results are observed for other LLMs in Appendix F.3. As for Weight-Activation Quantization, the quantized model does not show larger performance degradation in long texts than in short texts, as shown in Figure 7 (b).

For long-context tasks ($\geq 4k$), most LLMs are more sensitive to KV Cache Quantization than Weight-only and

Weight-Activation Quantization. With the same bit-width, the performance of LLMs with KV Cache Quantization is notably lower than that of Weight-only Quantization, as shown in Figure 7 (a, c). Within the LongChat (LLaMA-based) family, even KV8 quantization will cause notable performance degradation on long texts. For the Vicuna (LLaMA2-based) and ChatGLM family (32K), KV8 quantization is almost lossless, and the performance degradation on long texts occurs when using KV4. The Mistral family shows the highest tolerance to KV Cache Quantization, which is different from short texts. KV4 is still lossless for the Mistral family.

On LongEval and Multi-Doc QA benchmarks, within 2% accuracy loss, most LLMs can be quantized to W4, W4A8, or KV4 for short texts ($< 4k$), and W4, W4A8, or KV8 for long texts ($\geq 4k$), as illustrated in Table 1.

8. Limitations

In this paper, we focused solely on Post-training Quantization (PTQ) and did not consider Quantization-Aware Training (QAT). Furthermore, we did not conduct detailed ablation studies on certain hyperparameters, such as investigating the effect of different group sizes in group-wise quantization. For certain types of LLMs (e.g., small LLMs with sizes between 1B to 3B), we did not evaluate all the newly available LLMs. Some concrete recommendations on the bitwidth are specific to certain LLMs for certain tasks, and might not work well for new tasks and LLMs.

Acknowledgement

This work was supported by National Natural Science Foundation of China (No. 62325405, 62104128, U19B2019, U21B2031, 61832007, 62204164), Tsinghua EE Xilinx AI Research Fund, and Beijing National Research Center for Information Science and Technology (BNRist). We thank Boxun Li and Xiuhong Li from Infinigence-AI for their insightful suggestions and all anonymous reviewers for their valuable feedback.

Impact Statement

In this paper, we focus on the effects of quantization techniques on LLMs. Many of our findings can help the industry and academia reduce the large carbon footprint caused by LLM inference. However, as an evaluation paper, our extensive experiments also inevitably generated considerable carbon emissions. The broader social impacts of efficient machine learning techniques have already been widely discussed in other papers, none of which we feel must be specifically highlighted here.

References

- Almazrouei, E., Alobeidli, H., Alshamsi, A., Cappelli, A., Cojocar, R., Debbah, M., Goffinet, E., Heslow, D., Lounay, J., Malartic, Q., Noune, B., Pannier, B., and Penedo, G. Falcon-40B: an open large language model with state-of-the-art performance. 2023.
- Bisk, Y., Zellers, R., Gao, J., Choi, Y., et al. Piqa: Reasoning about physical commonsense in natural language. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pp. 7432–7439, 2020.
- Bondarenko, Y., Nagel, M., and Blankevoort, T. Quantizable transformers: Removing outliers by helping attention heads do nothing. *arXiv preprint arXiv:2306.12929*, 2023.
- Clark, P., Cowhey, I., Etzioni, O., Khot, T., Sabharwal, A., Schoenick, C., and Tafjord, O. Think you have solved question answering? try arc, the ai2 reasoning challenge. *arXiv preprint arXiv:1803.05457*, 2018.
- Cobbe, K., Kosaraju, V., Bavarian, M., Chen, M., Jun, H., Kaiser, L., Plappert, M., Tworek, J., Hilton, J., Nakano, R., Hesse, C., and Schulman, J. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*, 2021.
- Contributors, O. Opencompass: A universal evaluation platform for foundation models. <https://github.com/open-compass/opencompass>, 2023.
- Dettmers, T., Lewis, M., Belkada, Y., and Zettlemoyer, L. Llm.int8(): 8-bit matrix multiplication for transformers at scale. *arXiv preprint arXiv:2208.07339*, 2022.
- Du, Z., Qian, Y., Liu, X., Ding, M., Qiu, J., Yang, Z., and Tang, J. Glm: General language model pretraining with autoregressive blank infilling. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 320–335, 2022.
- Frantar, E., Ashkboos, S., Hoefler, T., and Alistarh, D. Gptq: Accurate post-training quantization for generative pre-trained transformers, 2023.
- Gao, L., Tow, J., Abbasi, B., Biderman, S., Black, S., DiPofi, A., Foster, C., Golding, L., Hsu, J., Le Noac’h, A., Li, H., McDonnell, K., Muennighoff, N., Ociepa, C., Phang, J., Reynolds, L., Schoelkopf, H., Skowron, A., Sutawika, L., Tang, E., Thite, A., Wang, B., Wang, K., and Zou, A. A framework for few-shot language model evaluation, 12 2023. URL <https://zenodo.org/records/10256836>.
- Gemma Team, T. M., Hardin, C., Dadashi, R., Bhupatiraju, S., Sifre, L., Rivière, M., Kale, M. S., Love, J., Tafti, P., Hussenot, L., and et al. Gemma. 2024. doi: 10.34740/KAGGLE/M/3301. URL <https://www.kaggle.com/m/3301>.
- Geva, M., Khashabi, D., Segal, E., Khot, T., Roth, D., and Berant, J. Did aristotle use a laptop? a question answering benchmark with implicit reasoning strategies. *Transactions of the Association for Computational Linguistics*, 9:346–361, 2021.
- GitHub. <https://github.com/features/copilot>. 2023.
- Gu, A. and Dao, T. Mamba: Linear-time sequence modeling with selective state spaces, 2023.
- Hendrycks, D., Burns, C., Basart, S., Critch, A., Li, J., Song, D., and Steinhardt, J. Aligning ai with shared human values. *Proceedings of the International Conference on Learning Representations (ICLR)*, 2021a.
- Hendrycks, D., Burns, C., Basart, S., Zou, A., Mazeika, M., Song, D., and Steinhardt, J. Measuring massive multitask language understanding. *Proceedings of the International Conference on Learning Representations (ICLR)*, 2021b.
- Huang, Y., Bai, Y., Zhu, Z., Zhang, J., Zhang, J., Su, T., Liu, J., Lv, C., Zhang, Y., Lei, J., Fu, Y., Sun, M., and He, J. C-eval: A multi-level multi-discipline chinese evaluation suite for foundation models. *arXiv preprint arXiv:2305.08322*, 2023.

- Jiang, A. Q., Sablayrolles, A., Mensch, A., Bamford, C., Chaplot, D. S., Casas, D. d. l., Bressand, F., Lengyel, G., Lample, G., Saulnier, L., et al. Mistral 7b. [arXiv preprint arXiv:2310.06825](#), 2023.
- Kadavath, S., Conerly, T., Askell, A., Henighan, T., Drain, D., Perez, E., Schiefer, N., Hatfield-Dodds, Z., DasSarma, N., Tran-Johnson, E., et al. Language models (mostly) know what they know. [arXiv preprint arXiv:2207.05221](#), 2022.
- Kim, S., Hooper, C., Gholami, A., Dong, Z., Li, X., Shen, S., Mahoney, M. W., and Keutzer, K. Squeezellm: Dense-and-sparse quantization. [arXiv preprint arXiv:2306.07629](#), 2023.
- Krishnamoorthi, R. Quantizing deep convolutional networks for efficient inference: A whitepaper. [arXiv preprint arXiv:1806.08342](#), 2018.
- Lai, G., Xie, Q., Liu, H., Yang, Y., and Hovy, E. Race: Large-scale reading comprehension dataset from examinations. [arXiv preprint arXiv:1704.04683](#), 2017.
- Lee, C., Jin, J., Kim, T., Kim, H., and Park, E. Owq: Lessons learned from activation outliers for weight quantization in large language models. [arXiv preprint arXiv:2306.02272](#), 2023.
- Levesque, H., Davis, E., and Morgenstern, L. The winograd schema challenge. In [Thirteenth international conference on the principles of knowledge representation and reasoning](#), 2012.
- Li, D., Shao, R., Xie, A., Sheng, Y., Zheng, L., Gonzalez, J., Stoica, I., Ma, X., and Zhang, H. How long can context length of open-source llms truly promise? In [NeurIPS 2023 Workshop on Instruction Tuning and Instruction Following](#), 2023.
- Lin, J., Tang, J., Tang, H., Yang, S., Dang, X., and Han, S. Awq: Activation-aware weight quantization for llm compression and acceleration, 2023.
- Lin, S., Hilton, J., and Evans, O. Truthfulqa: Measuring how models mimic human falsehoods, 2021.
- Liu, N. F., Lin, K., Hewitt, J., Paranjape, A., Bevilacqua, M., Petroni, F., and Liang, P. Lost in the middle: How language models use long contexts. [arXiv preprint arXiv:2307.03172](#), 2023a.
- Liu, P., Liu, Z., Gao, Z.-F., Gao, D., Zhao, W. X., Li, Y., Ding, B., and Wen, J.-R. Do emergent abilities exist in quantized large language models: An empirical study. [arXiv preprint arXiv:2307.08072](#), 2023b.
- Liu, Z., Oguz, B., Zhao, C., Chang, E., Stock, P., Mehdad, Y., Shi, Y., Krishnamoorthi, R., and Chandra, V. Llm-qat: Data-free quantization aware training for large language models. [arXiv preprint arXiv:2305.17888](#), 2023c.
- Liu, Z., Yuan, J., Jin, H., Zhong, S., Xu, Z., Braverman, V., Chen, B., and Hu, X. Kivi: A tuning-free asymmetric 2bit quantization for kv cache. [arXiv preprint arXiv:2402.02750](#), 2024.
- Mattern, J. and Hohr, K. Mamba-chat. GitHub, 2023. URL <https://github.com/havenhq/mamba-chat>.
- Nagel, M., Fournarakis, M., Amjad, R. A., Bondarenko, Y., Van Baalen, M., and Blankevoort, T. A white paper on neural network quantization. [arXiv preprint arXiv:2106.08295](#), 2021.
- OpenAI. Gpt-4 technical report, 2023.
- Paperno, D., Kruszewski, G., Lazaridou, A., Pham, Q. N., Bernardi, R., Pezzelle, S., Baroni, M., Boleda, G., and Fernández, R. The lambada dataset: Word prediction requiring a broad discourse context. [arXiv preprint arXiv:1606.06031](#), 2016.
- Park, G., Park, B., Kim, M., Lee, S., Kim, J., Kwon, B., Kwon, S. J., Kim, B., Lee, Y., and Lee, D. Lut-gemm: Quantized matrix multiplication based on luts for efficient inference in large-scale generative language models, 2023.
- Sakaguchi, K., Bras, R. L., Bhagavatula, C., and Choi, Y. Winogrande: An adversarial winograd schema challenge at scale. [Communications of the ACM](#), 64(9):99–106, 2021.
- Sap, M., Rashkin, H., Chen, D., LeBras, R., and Choi, Y. Socialiqa: Commonsense reasoning about social interactions. [arXiv preprint arXiv:1904.09728](#), 2019.
- Sheng, Y., Zheng, L., Yuan, B., Li, Z., Ryabinin, M., Fu, D. Y., Xie, Z., Chen, B., Barrett, C., Gonzalez, J. E., Liang, P., Ré, C., Stoica, I., and Zhang, C. Flexgen: High-throughput generative inference of large language models with a single gpu, 2023.
- Touvron, H., Martin, L., Stone, K., Albert, P., Almahairi, A., Babaei, Y., Bashlykov, N., Batra, S., Bhargava, P., Bhosale, S., et al. Llama 2: Open foundation and fine-tuned chat models. [arXiv preprint arXiv:2307.09288](#), 2023.
- Tow, J., Bellagente, M., Mahan, D., and Riquelme, C. Stablelm 3b 4elt. URL [<https://huggingface.co/stabilityai/stablelm-3b-4elt>] (<https://huggingface.co/stabilityai/stablelm-3b-4elt>).

- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., and Polosukhin, I. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
- Wan, Z., Wang, X., Liu, C., Alam, S., Zheng, Y., et al. Efficient large language models: A survey. *arXiv preprint arXiv:2312.03863*, 1, 2023.
- Wang, A., Singh, A., Michael, J., Hill, F., Levy, O., and Bowman, S. R. GLUE: A multi-task benchmark and analysis platform for natural language understanding. 2019. In the Proceedings of ICLR.
- Wang, B., Xu, C., Wang, S., Gan, Z., Cheng, Y., Gao, J., Awadallah, A. H., and Li, B. Adversarial glue: A multi-task benchmark for robustness evaluation of language models. *ArXiv*, abs/2111.02840, 2021.
- Wei, J., Bosma, M., Zhao, V., Guu, K., Yu, A. W., Lester, B., Du, N., Dai, A. M., and Le, Q. V. Finetuned language models are zero-shot learners. In *International Conference on Learning Representations*, 2022a.
- Wei, J., Tay, Y., Bommasani, R., Raffel, C., Zoph, B., Borgeaud, S., Yogatama, D., Bosma, M., Zhou, D., Metzler, D., et al. Emergent abilities of large language models. *arXiv preprint arXiv:2206.07682*, 2022b.
- Wei, X., Zhang, Y., Zhang, X., Gong, R., Zhang, S., Zhang, Q., Yu, F., and Liu, X. Outlier suppression: Pushing the limit of low-bit transformer language models. *Advances in Neural Information Processing Systems*, 35:17402–17414, 2022c.
- Workshop, B., Scao, T. L., Fan, A., Akiki, C., Pavlick, E., Ilić, S., Hesslow, D., Castagné, R., Luccioni, A. S., Yvon, F., et al. Bloom: A 176b-parameter open-access multilingual language model. *arXiv preprint arXiv:2211.05100*, 2022.
- Xiao, G., Lin, J., Seznec, M., Wu, H., Demouth, J., and Han, S. Smoothquant: Accurate and efficient post-training quantization for large language models. In *International Conference on Machine Learning*, pp. 38087–38099. PMLR, 2023.
- Yao, Z., Yazdani Aminabadi, R., Zhang, M., Wu, X., Li, C., and He, Y. Zeroquant: Efficient and affordable post-training quantization for large-scale transformers. *Advances in Neural Information Processing Systems*, 35: 27168–27183, 2022.
- Yao, Z., Wu, X., Li, C., Youn, S., and He, Y. Zeroquant-v2: Exploring post-training quantization in llms from comprehensive study to low rank compensation. *arXiv preprint arXiv:2303.08302*, 2023.
- Yuan, T., Ning, X., Zhou, D., Yang, Z., Li, S., Zhuang, M., Tan, Z., Yao, Z., Lin, D., Li, B., et al. Lv-eval: A balanced long-context benchmark with 5 length levels up to 256k. *arXiv preprint arXiv:2402.05136*, 2024.
- Yuan, Z., Niu, L., Liu, J., Liu, W., Wang, X., Shang, Y., Sun, G., Wu, Q., Wu, J., and Wu, B. Rptq: Reorder-based post-training quantization for large language models. *arXiv preprint arXiv:2304.01089*, 2023.
- Zellers, R., Holtzman, A., Bisk, Y., Farhadi, A., and Choi, Y. Hellaswag: Can a machine really finish your sentence? *arXiv preprint arXiv:1905.07830*, 2019.
- Zhang, S., Roller, S., Goyal, N., Artetxe, M., Chen, M., Chen, S., Dewan, C., Diab, M., Li, X., Lin, X. V., et al. Opt: Open pre-trained transformer language models. *arXiv preprint arXiv:2205.01068*, 2022.
- Zheng, C., Huang, M., and Sun, A. Chid: A large-scale chinese idiom dataset for cloze test. *arXiv preprint arXiv:1906.01265*, 2019.
- Zheng, L., Chiang, W.-L., Sheng, Y., Zhuang, S., Wu, Z., Zhuang, Y., Lin, Z., Li, Z., Li, D., Xing, E., et al. Judging llm-as-a-judge with mt-bench and chatbot arena. *arXiv preprint arXiv:2306.05685*, 2023a.
- Zheng, L., Chiang, W.-L., Sheng, Y., Zhuang, S., Wu, Z., Zhuang, Y., Lin, Z., Li, Z., Li, D., Xing, E., et al. Judging llm-as-a-judge with mt-bench and chatbot arena. *arXiv preprint arXiv:2306.05685*, 2023b.
- Zhou, Z., Ning, X., Hong, K., Fu, T., Xu, J., Li, S., Lou, Y., Wang, L., Yuan, Z., Li, X., et al. A survey on efficient inference for large language models. *arXiv preprint arXiv:2404.14294*, 2024.

Appendix

A. Additional Preliminaries

A.1. Large Language Model Inference

As mentioned in Sec. 1, the transformer-based (Vaswani et al., 2017) LLMs have two distinctive stages, including the prefill and the decoding stages. Take batch size = 1 as an example. **During the prefill stage**, a prompt sequence is utilized to generate the next token. Simultaneously, the Key and Value tensors of each transformer block in LLM are generated and stored as the KV Cache. The primary operator in the prefill stage is the General Matrix Multiply (GEMM). **In the decoding stage**, the LLM takes one generated token from step t as the input and uses the KV Cache to generate the next token of step $t + 1$. The generation of the current token depends on one previously generated token and the KV Cache. The main operator in the decoding stage is the General Matrix-Vector Multiply (GEMV). Furthermore, as the length of the context increases, the memory overhead introduced by the KV Cache linearly grows.

A.2. Quantization

As illustrated in Sec. 2.1, we use symmetric uniform quantization for Activations and group-wise asymmetric uniform quantization for the Weights and KV Cache. The quantization formats we chose are commonly used in existing quantization methods. For example, LLM.int8() (Dettmers et al., 2022) and SmoothQuant (Xiao et al., 2023) use symmetric per-token quantization for activation, AWQ (Lin et al., 2023) and GPTQ (Frantar et al., 2023) asymmetric group-wise quantization for weight, and Flexgen (Sheng et al., 2023) and KIVI (Liu et al., 2024) use asymmetric group-wise quantization for KV Cache. Specifically, the reason for our design choice of the quantization format is shown below:

1. To reduce the memory overhead, which is extremely important for LLM inference, many existing methods quantize weights and KV Cache to a very low bit-width (≤ 4 -bit). On the one hand, to maintain accuracy, existing methods typically employ fine-grained asymmetric group-wise quantization. On the other hand, during inference, low-bit-width data is de-quantized to higher-bit-width data, such as FP16, for computation. The overhead of de-quantizing a fine-grained group-wise formatted tensor and a coarse-grained tensor is quite similar. Therefore, using a fine-grained group-wise format for weight and KV cache is appropriate.
2. To reduce computational costs, existing work quantizes both Activation and Weight to lower precision, enabling the use of faster low-precision computing units. If asymmetric group-wise quantization is applied to activation, the low-precision computing units cannot be utilized directly. A more fine-grained quantization format for activation needs to come with a complex GPU kernel design, which will involve a longer development time and could only bring a small improvement [7]. Therefore, for efficient and straightforward utilization of low-precision computing units, symmetric per-token quantization is typically used for activation.

In addition, in this paper, we only focus on Weight-only, Weight-Activation, and KV Cache quantization, without Activation-only quantization. The reason is that we aim to use quantization to reduce computation and memory overhead in the inference process. As mentioned in Sec. 1, the prefill stage and decoding stage are mainly compute-bounded and memory-bounded, respectively. Furthermore, the memory requirements for KV Cache increase linearly as the context length grows larger. To address the above three challenges, we can employ Weight-only, Weight-Activation, and KV Cache quantization.

In contrast, Activation-only quantization makes it difficult to accelerate the inference process of LLMs. For the prefill stage, since the weights are not quantized, we can not use low-precision computing units to accelerate computations. In the decoding stage (with a small batch size), the bottleneck that limits inference speed lies in slow weight access, and saving activation memory does not lead to notable acceleration.

It’s also worth mentioning that Activation-only quantization, while difficult to speed up the inference process, can significantly reduce communication and memory costs during training. As our work focuses on accelerating the inference process of LLMs, we didn’t consider activation-only quantization, but we will include this discussion in our paper to highlight it as a design choice to be carefully considered in training optimization.

A.3. Experimental Setup Details

For the group-wise KV Cache and Weight-only Quantization, we set the group size to be the hidden dimension size of one head in the model’s multi-head attention block. Specifically, for the Mistral, LLaMA2, Vicuna, LongChat, and ChatGLM families, the group size is 128. For the Falcon family, the group size is 64. For the Bloomz and OPT families, different LLMs have different group sizes. The OPT-6.7B, OPT-13B, OPT-30B, OPT-66B, Bloomz-1B7, Bloomz-7B1, and Bloomz-175B have the same group size of 128. The OPT-125M, OPT-1.3B, and Bloomz-560M have the same group size of 64. The OPT-2.7B and Bloomz-3B have the same group size of 80. Finally, only Bloomz-1B1 has a group size of 96. For StableLM-3B, we use a group size of 80.

B. Additional Details of Evaluation on Basic NLP Abilities

B.1. Introduction of Datasets

We evaluate the basic NLP tasks of the quantized LLMs in three aspects, namely language modeling, understanding, and reasoning. Each aspect of the abilities is measured based on the performances of quantized LLMs on certain datasets. We evaluate the language modeling ability on the CHID (Zheng et al., 2019) and Winogrande (Sakaguchi et al., 2021) datasets, understanding ability on the RACE (Lai et al., 2017) and LAMBADA (Paperno et al., 2016) datasets, and reasoning on the PIQA (Bisk et al., 2020) and SIQA (Sap et al., 2019) datasets for their frequent usage. All the evaluation experiments are conducted based on the datasets integrated by the opencompass project (Contributors, 2023). Below is a brief introduction to the related datasets.

The CHID dataset (Zheng et al., 2019) is a Chinese idiom reading comprehension task, which requires the evaluated LLMs to select the correct idiom to fill in the blank according to the context, with several candidate idioms. The dataset is split into the train/dev/test sets. We evaluate the quantized LLMs on the test split, containing 2,002 test sentences. Human performance serves as an upper bound, which is 87.1% on the test split.

The Winogrande dataset (Sakaguchi et al., 2021) is a large-scale dataset of 44k problems, inspired by the Winograd Schema Challenge (Levesque et al., 2012), with both the scale and the hardness improved. Each question of the dataset is composed of a sentence and a pronoun. The evaluated LLM is required to judge what the pronoun refers to according to the context. The whole dataset is divided into train/dev/test sets. The evaluations are based on the dev set, with 1,267 test sentences involved. Human accuracy on the dev set is 94.1%.

The RACE dataset (Lai et al., 2017) is a large-scale reading comprehension dataset with over 28,000 passages and nearly 100,000 questions. The dataset is collected from English examinations in China, designed for middle and high school students. The dataset is split into train/dev/test sets and the evaluations take place in the test set, which contains 3,498 questions. We report model performances on the high school part of the dataset because it’s harder and has more test instances than the middle school part. Human accuracy on the high school questions is 94.2%.

The LAMBADA dataset (Paperno et al., 2016) evaluates the capabilities of the LLMs for text understanding through a word prediction task. LAMBADA is a collection of narrative passages sharing the characteristic that human subjects can guess the last word if they have access to the whole passage but not if they only see the last sentence preceding the target word. The LAMBADA dataset is extracted from BookCorpus and consists of 10,022 passages, divided into 4,869 development and 5,153 test passages. We evaluate the quantized LLMs using the test set.

The PIQA dataset (Bisk et al., 2020) is a physical interaction question-answering task designed to test the model’s knowledge of physical commonsense, which requires the models to choose the most reasonable solution based on the given scenario and two possible solutions. This dataset consists of 16k training samples, 800 development samples, and 2k test samples. We conduct evaluations on the test split, where a 95% human accuracy serves as an upper bound.

The SIQA dataset (Sap et al., 2019) is a social interaction question-answering task designed to test the model’s knowledge of social commonsense, which requires the models to choose the most reasonable behavior given a scenario and three possible subsequent behaviors. This dataset contains 38,963 training samples, 1,951 development samples, and 1,960 test samples. We benchmark the quantized LLMs on the development samples. Human performance on the development split is 86.9%.

B.2. Introduction of Metrics

We employ two basic methods, “Evaluating in the PPL mode”, and “Evaluating in the Gen (short for generation) mode”, to evaluate the performance of the quantized LLMs.

The first method, “**Evaluating in the PPL mode**”, is designed to evaluate multiple-choice tasks. We first combine the original questions and each choice into several narrative sentences. LLMs will calculate the perplexity (PPL) of each narrative sentence. The corresponding choice of the narrative sentence with the lowest PPL scores is the model’s answer. If the model’s answer matches the ground-truth answer, it is deemed correct; otherwise, it is considered incorrect. Below is an example from the Winogrande dataset. In this example, we need to choose the most appropriate word from two options to fill in the blank space in the original question. Therefore, we respectively fill the two words into the original question to create two narrative sentences for evaluation.

An example of “Evaluating in the PPL Mode” from Winogrande dataset (Question 37)

Original question: Joel researched laws and helped to open a preschool for Eric. Because ___ is very good with kids.

Option 1: Joel

Option 2: Eric

Reorganized narrative sentence 1: Good sentence: Joel researched laws and helped to open a preschool for Eric. Because Joel is very good with kids.

Reorganized narrative sentence 2: Good sentence: Joel researched laws and helped to open a preschool for Eric. Because Eric is very good with kids.

The second method, “**Evaluating in the Gen mode**”, will simply combine the original question and options (if any) into a single prompt. At the same time, an instruction is added to each prompt to guide LLMs in generating the correct answer. The LLMs under evaluation are required to generate replies according to the prompt, which will be post-processed to get the final answer. Below is a sample extracted from the LAMBADA dataset. In this example, we don’t have any options, so we combine the instruction “Please complete the following sentence:” and the original question as the final prompt. For the LAMBADA dataset, we only extract the next word of the original question from the generated text as the answer. If the answer matches the ground-truth answer, it is deemed correct; otherwise, it is considered incorrect. For other tasks, there may be different post-processing methods for the final answer.

An example of “Evaluating in the Gen Mode” from LAMBADA dataset (Question 27)

Original question: She kisses me again. “Oh-kay!” Jen whistles. “Can we go now? I don’t wanna break up the lovey-dovey show goin’ on, but I have shopping to do.” I laugh and release Lexy, wrapping a friendly arm around Jen’s shoulders. “Oh, how dull my life would be without you,

Reorganized question as prompt: Please complete the following sentence: She kisses me again. “Oh-kay!” Jen whistles. “Can we go now? I don’t wanna break up the lovey-dovey show goin’ on, but I have shopping to do.” I laugh and release Lexy, wrapping a friendly arm around Jen’s shoulders. “Oh, how dull my life would be without you,

Specifically, we adopt the “evaluating in the ppl mode” method to evaluate the LLMs on the CHID, Winogrande, RACE, PIQA, and SIQA datasets, which are originally organized in the choice question format. As for the generative task LAMBADA, we adopt the “evaluating in the gen mode” method. Among all LLMs, the OPT family has a notably terrible performance on the CHID and RACE datasets, which is quite close to random guesses. Hence, we rule out the model-dataset combinations in the following discussion.

In addition, for each LLM, we plot the performance curves under different bit-width by averaging the normalized performance on different datasets, as shown from Figure 8 to Figure 12.




Figure 8: The effect of quantization on the OPT family




Figure 9: The effect of quantization on the Bloomz family




Figure 10: The effect of quantization on the Falcon family




Figure 11: The effect of quantization on the LLaMA2 family




Figure 12: The effect of quantization on the Mistral family

B.3. Additional Results on Different Tensor Types

For Weight-only Quantization, the larger the model size, the higher the tolerance for Weight-only Quantization. Besides, most models can perform nearly as well as their FP16 counterparts when quantizing to W4. Most LLMs suffer from slight accuracy loss at W3 quantization. The OPT family is an exception in that most OPT models have severe accuracy loss when quantizing to only W3, as shown in Figure 8 (a). Most LLMs collapse at W2 quantization except the largest Bloomz-175B and Falcon-180B, as shown in Figure 9 (a) and Figure 10 (a).

For Weight-Activation Quantization, the larger the model size, the lower the tolerance for Weight-Activation Quantization, which is opposite from Weight-only Quantization. Almost all the families work well under the W8A8 and W4A8 quantization. OPT model family is also an exception in which even W8A8 can cause significant performance loss to the OPT-66B model for all tasks. For models that are larger than 6.7B, we cannot quantize them to W4A8 on all tasks, as shown in Figure 8 (b). For models that are smaller than 6.7B, we can quantize them to W4A8 without significant accuracy, which is similar to other LLMs. In our experiments, W4A4 quantization results in significant accuracy loss for all LLMs, with only two exceptions: Bloomz-560B and Bloomz-1B1, as shown in Figure 9 (b). In addition, when quantizing most LLMs to W4A4, the largest model from the same model family usually has the lowest performance. We only observe a few exceptions, such as the evaluation of the Bloomz family on the Winogrande task, where the worst model is Bloomz-3B instead of Bloomz-175B.

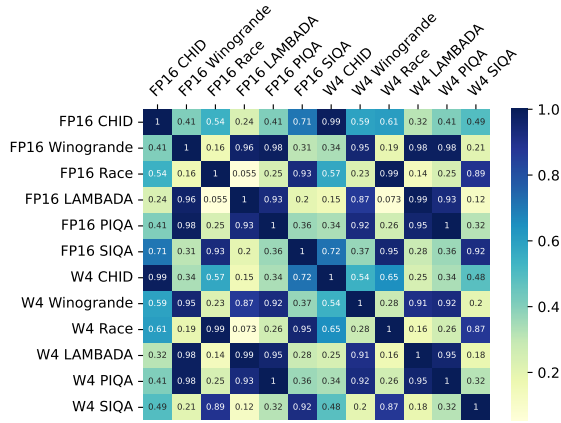
For KV Cache Quantization, the larger the model size, the higher the tolerance for KV Cache Quantization, which is similar to Weight-only Quantization. Nearly all the LLM families undergo nearly no accuracy loss when quantizing to KV4. Most LLMs have a slight accuracy loss at KV3 quantization. Some good cases can be found in the Falcon family that the larger Falcon-40B and Falcon-180B nearly have no accuracy loss on LAMBADA, PIQA, and SIQA datasets. Similar results also appear on the largest LLaMA2-70B from the LLaMA2 family on LAMBADA, PIQA, and SIQA datasets. Some bad cases can also be observed in our results, such as the LLaMA2-13B from the RACE dataset, that the larger LLaMA2-13B model has a larger performance loss than LLaMA2-7B. Besides, the accuracy loss of KV Cache Quantization is usually less than that of W3 quantization, especially for the larger LLMs, in which the KV3 quantized models have only slight accuracy loss. However, the performance of the W3 quantized smaller models has a larger accuracy loss, especially OPT-2.7B and OPT-66B, as shown in Figure 8 (a, c).

Specifically, for many real-world applications, using **both the Weight-only and KV Cache Quantization** is necessary to alleviate the large memory overhead introduced by the model size (Weight), large batch sizes (KV Cache), and long texts (KV Cache). As shown in Table 7, we apply both Weight-only Quantization and KV Cache Quantization LLaMA2 family on the LAMBADA dataset to evaluate the basic understanding ability, MT-bench to evaluate dialogue ability, and LongEval dataset to evaluate long-context processing ability. For LAMBADA and MT-bench, in most cases, quantizing LLMs to W4KV4 has only slight performance loss (< 2%), and W8KV4 is a better choice for all the tested LLMs. For LongEval, LLMs are more sensitive to KV Cache Quantization, as discussed in Sec. 7. W8KV4 can cause significant accuracy loss by over 10%. W4A8 is much better for the Vicuna family, and W8A8 is nearly lossless.

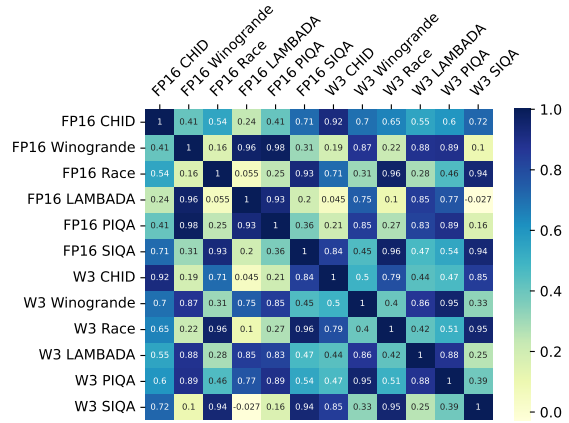
B.4. Additional Results on Different LLMs

The OPT model family is more sensitive to both the Weight-only and the Weight-Activation Quantization. For the Weight-only Quantization, the OPT family exhibits higher Kurtosis in Weight tensors compared to other LLMs, leading to the collapse of most OPT models when quantized to W3 instead of W2. In the case of Weight-Activation Quantization, the

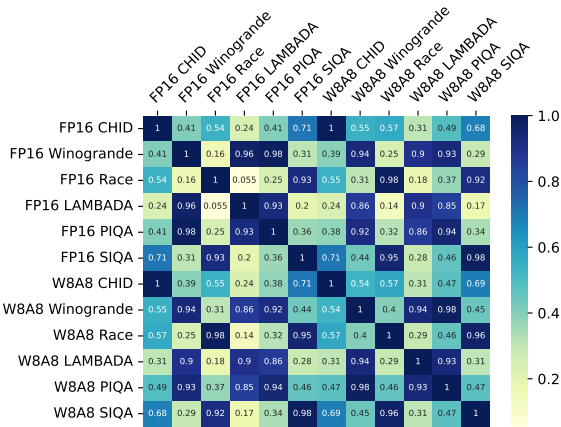
Evaluating Quantized Large Language Models



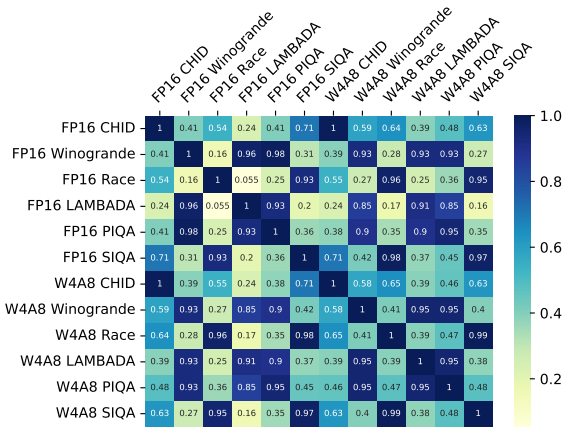
(a) W4 Quantization




(b) W3 Quantization




(c) W8A8 Quantization



(d) W4A8 Quantization



(e) KV4 Quantization



(f) KV3 Quantization

Figure 13: Spearman correlation between each pair of tasks and quantization.

Evaluating Quantized Large Language Models

Table 7: The effects of applying both Weight-only and KV Cache Quantization. For LAMBADA, we use the base LLaMA2 models. For MT-Bench, we use LLaMA2-Chat models. For LongEval, we use the Vicuna-7B and Vicuna-13B models based on the LLaMA2 family.

LLaMA2	LAMBADA				MT-Bench (Avg.)				LongEval (16K)			
	FP16	W8KV8	W8KV4	W4KV4	FP16	W8KV8	W8KV4	W4KV4	FP16	W8KV8	W4KV8	W8KV4
7B	73.32	73.37	72.07	70.54	4.73	4.72	4.76	4.68	57.80	56.40	59.60	37.00
13B	76.48	-	-	75.02	5.39	5.6	5.25	5.28	41.60	40.80	36.40	29.00
70B	78.96	-	-	77.72	6.13	6.01	5.98	5.64	-	-	-	-

performance of most OPT models collapses at W4A8 and even at W8A8 instead of W4A4. As indicated in Table 3, this phenomenon is attributed to the higher Kurtosis of Activation in the OPT models compared to other LLMs. Specifically, OPT-6.7B has an average Kurtosis of 1562.67, surpassing that of LLaMA2-70B (1279.15). Consequently, within the OPT family, only OPT models smaller than 6.7B can be quantized into W4A8. Moreover, OPT-66B exhibits a significantly higher average Kurtosis of Activation (4945.32) compared to OPT-6.7B, making OPT-66B unable to be quantized to even W8A8.

As discussed in Sec. 3.3, we mentioned that for the majority of models, the performance order of the Weight-only and KV Cache quantized models is generally consistent with that of the FP16 models.

We further investigate the performance correlation between the quantized LLMs and the original FP16 LLMs. We select two representative models from each model family (typically a small model together with a large model in the same family) to form a model set. Specifically, we select the LLaMA2-7B and LLaMA2-70B from the LLaMA2 family, the Falcon-7B, and Falcon-180B from the Falcon family, and the Bloomz-7B1 and Bloomz-175B from the Bloomz family and the Mistral-7B and Mixtral-8x7B from the Mistral family. We also select the CHID, Winogrande, RACE, LAMBADA, PIQA, and SIQA datasets introduced in Appendix B.1. For each dataset, both the quantized LLMs and the original FP16 LLMs has a certain performance order. For example, on the CHID dataset, we have LLaMA2-70B > LLaMA2-7B > ... To evaluate the correlation between the quantized LLMs and the original LLMs, we calculate the Spearman Correlation of Quantized LLM and FP16 LLMs on different datasets. The Spearman Correlation can be calculated by the following formula:

$$\rho = 1 - \frac{6 \sum d_i^2}{n(n^2 - 1)}, \quad (3)$$

where ρ represents the Spearman Correlation coefficient, d_i represents the differences in ranks of variables, and n is the number of samples. In addition, the Spearman correlation coefficient ρ ranges from -1 to 1. $\rho = 1$ indicates a perfect positive monotonic relationship, $\rho = -1$ indicates a perfect negative monotonic relationship, and $\rho = 0$ indicates no monotonic relationship.

The Spearman correlation results are shown in Figure 13. In this Figure, the top-left 6×6 square represents the Spearman correlation of the FP16 LLMs on different datasets. The top-right and the bottom-left 6×6 square represent the Spearman correlation of the FP16 LLMs and the quantized LLMs on different datasets. Finally, the bottom-right 6×6 square represents the Spearman correlation of the quantized LLMs on different datasets. Based on these results, we can draw the following two observations: First, the performance order is highly task-relevant, which may be attributed to the diverse training settings and data each model trained with. Secondly, we observe that Weight-only, Weight-Activation, and KV Cache quantization have a distinct order-preserving character that the Spearman correlation coefficients between the quantized LLMs and their FP16 counterparts on the same task are quite high (more than 0.9, typically) when the bit-width is higher than W4, W4A8, and KV4.

In Sec 3.3, we discover that the performance of the Mixtral-8x7B is close to LLaMA2-70B, and the sensitivity of Mixtral-8x7B to quantization is closer to that of the smaller LLaMA2-7B and even higher than that of the Mistral-7B from the same model family as shown in Figure 12. Similar results can also be found on different datasets and different LLM families. Specifically, we find that the Mixtral-8x7B MoE model is always more sensitive to quantization than Falcon-40B and Falcon-180B. In addition, to evaluate the effect of quantization on the gate layer in MoE models, we keep the gate layer in the Mixtral-8x7B as FP16 while quantizing other linear layers. Interestingly, we find no accuracy gain on the LAMBADA dataset, which means that whether the gate layer is quantized has little impact on the final performance of the model. As a promising method to increase the efficiency of LLMs, MoE technology has attracted widespread attention. In the future, how to simultaneously leverage quantization methods and MoE techniques to enhance inference efficiency is a direction that

Table 8: The PPL statistical results for the W8A8 Mistral-7B model on multiple-choice task RACE.

Bit-width	C1		C2		C3		C4	
	PPL	Std	PPL	Std	PPL	Std	PPL	Std
FP16	1.996	0.0016	2.020	0.0011	2.021	0.0013	2.021	0.0011
W8A8	2.011	0.0015	2.035	0.0010	2.036	0.0012	2.035	0.0010

requires further in-depth research.

B.5. Additional Results on Different Tasks

Multi-choice accuracy sometimes cannot accurately reflect the performance. For multiple-choice tasks, we concatenate each answer candidate with the question to form multiple sentences (see Appendix B.2). Then, we compute the perplexity (PPL) of each sentence and choose the sentence with the lowest PPL as the answer. We observe that in some cases, quantization even brings notable accuracy gain. For example, the accuracy of Mistral-7B on the RACE dataset increases by 3.5% after W8A8 quantization. We look into this unexpected phenomenon as follows. We categorize multiple-choice questions into four classes, as presented in Table 8. C1 and C4 denote questions for which both the FP16 and W3 Mistral-7B answer correctly and incorrectly, respectively. C2 denotes questions answered correctly before quantization but incorrectly after quantization, while C3 is the opposite. We compute the average PPL for the correct answers and the average standard deviation (Std) of PPL within each of the four categories. We observe that (1) The average PPL increases after quantization, indicating a decline in performance for quantized LLMs. (2) The average Std in C2 and C3 is notably smaller than that in C1 and C4. This suggests that questions in C2 and C3 are instances where the model exhibits high uncertainty. Consequently, it is possible for the quantized model to guess the correct answer, resulting in the phenomenon that quantization brings accuracy gain.

We observe many cases in which the quantized LLMs have better performance compared to the FP16 counterparts; some examples are listed as follows. For the OPT family, the W4 OPT-30B has a slight accuracy gain on both Winogrande and PIQA tasks. For the Mistral family, we observe that after Weight-only and Weight-Activation Quantization, the Mistral-7B and Mixtral-8x7B have a significant accuracy gain on the RACE and SIQA datasets. For the ChatGLM family, we find that the W3 and W4A8 quantized ChatGLM3-6B-32K models have significant improvements on the RACE dataset. The ChatGLM3-6B model, which is designed for short texts, also has performance gains on RACE and SIQA datasets. For the Bloom family, the Bloomz-560, Bloomz-1B7, and Bloomz-3B have significant accuracy gain when quantized to W3, W4A8, and KV3. The reason has already been discussed in Sec. 3.4. However, in some cases, the quantized LLM does not consistently show a higher PPL. As shown in Table 9, we summarize the average PPL of right answers and average Std of each of four options for OPT-30B on the Winogrande dataset and categorize multiple-choice questions into four classes. As described in Sec. 3.4, C1 and C4 represent questions on which both the FP16 OPT-30B and the W4 OPT-30B answer correctly and incorrectly, respectively. C2 represents questions that were answered correctly before quantization but incorrectly after quantization, while C3 is the opposite. We find different results that the PPL of the questions in C1, C3, and C4 slightly decreases after quantization, which means the quantized LLMs are not significantly getting worse. The Std of the PPL for each option in C2 and C3 is still significantly smaller than that in C1 and C4, which is similar to the analysis in Sec. 3.4. The quantized model has lower accuracy because it does not guess uncertain questions correctly. This result further supports our conclusion that tasks in the form of multiple-choice questions sometimes cannot accurately reflect the performance of LLMs.

Table 9: The PPL statistical results of the W4 OPT-30B model on multiple-choice task Winogrande.

Bit-width	C1		C2		C3		C4	
	PPL	Std	PPL	Std	PPL	Std	PPL	Std
FP16	4.06	0.054	4.23	0.010	4.24	0.009	4.18	0.041
W3	4.04	0.054	4.24	0.009	4.20	0.010	4.16	0.042

B.6. Additional Results on Different Quantization methods

Restoring performance to a level similar to the FP16 LLMs is challenging with SOTA quantization methods when using extremely low bit-width, such as W2 or W4A4. As shown in Table 5, we observe that for the W3 quantization,

Evaluating Quantized Large Language Models

Table 10: The statistical results on different linear types in the OPT family.

Model	W or A	Q		K		V		O		FC1		FC2	
		Kuro.	Std	Kuro.	Std	Kuro.	Std	Kuro.	Std	Kuro.	Std	Kuro.	Std
OPT-1.3B	W	4.23	0.02	6.14	0.02	3.73	0.02	55.69	0.02	4.13	0.02	5.04	0.02
	A	494.20	1.00	494.20	1.00	494.20	1.00	99.04	0.20	467.87	1.01	1214.35	0.15
OPT-6.7B	W	3.82	0.02	4.40	0.02	3.29	0.01	33.80	0.01	3.51	0.02	3.66	0.02
	A	1781.86	1.00	1781.86	1.00	1781.86	1.00	62.23	0.20	1733.60	1.01	2234.63	0.15
OPT-66B	W	4.02	0.01	4.73	0.01	3.25	0.01	11.87	0.01	3.27	0.01	4.01	0.01
	A	5007.77	1.00	5007.77	1.00	5007.77	1.00	177.45	0.14	4783.57	1.01	9687.57	0.11

Table 11: The statistical results on different linear types in the Mamba family.

Model	W or A	X_Proj		Out_proj	
		Kuro.	Std	Kuro.	Std
Mamba-2.8B	W	5.10	0.04	3.45	0.02
	A	871.29	0.18	26836.54	0.22

AWQ can improve the performance of quantized LLMs. Moreover, the larger the model, the closer the performance is to the FP16 baseline. However, in the case of W2 quantization, where quantized LLMs lose their abilities entirely, AWQ cannot restore the corrupted performances. Similar results are also observed with Weight-Activation Quantization. While SmoothQuant can partially recover the performance for W4A4 quantization, the performance remains significantly lower than the FP16 baseline. The conclusions are applicable to other tasks as well, and additional details can be found in our Appendix. Specifically, the results for emergent abilities are shown in Table 15, dialogue results can be found in Table 19, and long-context evaluations are depicted in Figure 32.

B.7. Addition Results on Statistical Analysis

We show the insights of two directions introduced in Sec. 3.2 based on the statistical analysis.

The insights of Mix-precision Quantization: We discover that different tensor type has very different data distributions. As shown in Table 10, the Weight tensors of output projection layers have the largest kurtosis among different linear layers. However, the Activation tensors of output projection layers have the smallest kurtosis among different linear layers. This phenomenon means that the Weight of the output projection layers may need higher bit-width, and the Activation tensors of the output projection layers may need lower bit-width than other linear layers. The FC2 layers have the largest kurtosis in Activation tensors, which may indicate that a higher bit-width is needed. Besides, the Q, K, V, and FC1 have similar kurtosis in both the Weight and Activation tensors, which may need the same bit-width. Similar phenomena also appear in the LLaMA2 family. As shown in Table 4, the most significant phenomenon is that the Activation tensors of the down projection layers have the largest kurtosis, which is significantly larger than that of other linear layers. From these phenomena, we recommend more studies on the mix-precision quantization methods to get a better trade-off between hardware efficiency and performance.

The insights of Different Scaling Schemes: Different input data can generate different Activation tensors, while Weights are shared across all data. Therefore, in most cases, we pre-quantize the Weight tensors, eliminating the need to quantize Weight tensors during each inference. For Activation Quantization, we mainly have two choices: (1) We can calculate the scaling factors offline based on the Activation tensors of some calibration samples, which is called **static quantization**. During inference, we can directly use the pre-defined scaling factors to quantize Activation tensors. (2) We can use the runtime Activation tensors to calculate scaling factors online, which is called **dynamic quantization**. In this case, we need additional processing to calculate scaling factors for each token, which will cause additional computation overhead. We evaluate the LLaMA2 family with both static and dynamic quantization on the Wikitext dataset and report the PPL (the lower, the better). As shown in Table 12, for the W4A8 LLaMA2 model, the performance loss caused by static quantization is much greater than dynamic quantization, especially for the LLaMA2-7B and LLaMA2-70B models. According to the

statistical information in Table 4, there are significantly more outliers in the down projection layers than in other layers, which may be the reason for the significant loss caused by static quantization. Therefore, we apply dynamic quantization to all down projection layers and static quantization to the remaining layers in each LLaMA2 model. This strategy results in a significant improvement in the PPL, which is very close to the dynamic quantized LLaMA2 models.

Table 12: The performance of the static quantization and dynamic quantization for Activation tensors on Wikitext dataset.

Models	FP16	W4A8 (Dynamic)	W4A8 (Static)	W4A8 (Static w.o. Down)
LLaMA2-7B	11.71	12.51	40.78	12.63
LLaMA2-13B	10.22	10.59	12.10	10.57
LLaMA2-70B	6.87	7.23	16.89	7.26

C. Additional Details of Evaluation on Emergent Abilities

C.1. Introduction of Datasets

The MMLU dataset is a comprehensive dataset including 57 tasks spanning across disciplines such as mathematics, computer science, history, etc. All the questions are presented as multiple-choice questions and are derived from a variety of educational assessments, including exam questions from courses and exercises in books. In total, the dataset consists of 15,908 questions, split into the “dev subset”, the “validation subset”, and the test set. The “dev subset” has five questions per task and is designed for few-shot prompt evaluation. The “validation subset” has 1540 questions, and the test set contains 14079 questions.

The CEval dataset is a Chinese dataset consisting of 13,948 multiple-choice questions distributed across 52 academic subjects, ranging from humanity to engineering. These questions are further categorized into four difficulty levels, spanning from middle school to professional tests, and are sourced from mock exams available on the Internet.

The ARC dataset comprises a range of multiple-choice questions whose target is students from grades three to nine. It is divided into two subsets, namely the easy set and the challenge set, based on the performance of two baseline solvers. With a total of 7,787 questions, the dataset includes 5,197 questions in the easy set and 2,590 questions in the more demanding challenge set.

The Hellaswag dataset focuses on commonsense reasoning through a collection of multiple-choice questions. Drawing from two distinct sources, namely the ActivityNet Captions dataset and WikiHow, the dataset is made up of a total of 70k questions. Specifically, 25k questions are sourced from the ActivityNet Captions dataset, while the remaining 45k questions are from WikiHow, both representing the highest human agreement.

The StrategyQA dataset contains 2,290 multi-step reasoning questions in its train set, each accompanied by a definite yes/no answer.

The GSM8K dataset comprises a total of 8,500 instances, with 7,500 training problems and an additional 1,000 test problems. The dataset focuses on grade school mathematics problems that can be solved through two to eight arithmetic steps. Human workers craft all the problems without linguistic templates. Solutions are expressed in natural language.

C.2. Introduction of Metrics

To evaluate the In-Context Learning ability of LLMs on the MMLU and CEval datasets, we employ the PPL mode described in Appendix B.2. Specifically, for the five-shot experiments, few examples are extracted from the “dev subset” of the datasets. In the case of ten-shot experiments, given the limited quantity of data in the “dev subset”, examples are drawn from both the “dev subset” and the “validation subset”. This approach ensures a comprehensive evaluation by avoiding any duplication between examples and questions.

We evaluate the Instruction-Following ability of LLMs by employing the templates provided by Wei et al. (2022a). These templates serve as explicit instructions guiding the language model in completing the tasks within the dataset. Questions in one dataset are embedded in a fixed template, requiring the LLMs to respond accordingly.

An example of Template Utilization in Hellaswag dataset. Templates are highlighted in bold.

What most naturally follows?

You must be at least 18 years of age.

- A. You must have no felony convictions or drug offenses. You must have a two-year college education.
- B. If you are 16 or 17, you'll need your parents' or guardians' consent. Anyone over 18 who is under a guardian's care must have their guardian's permission.
- C. You must have a driver's license. You must be at least 21 years of age.
- D. You must have any past felony convictions involving marijuana. You must be able to work full time.

An example of Template Utilization in ARC dataset. Templates are highlighted in bold.

Question: What does photosynthesis produce that helps plants grow?

What is the correct answer to the question from the following choices?

- A. water
- B. oxygen
- C. protein
- D. sugar

To evaluate the Multi-Step Reasoning ability of LLMs, we employ the GEN mode, as described in Appendix B.2. The LLMs generate predictions by considering the provided questions along with a few chain-of-thought examples. For StrategyQA, there are six examples, and for GSM8K, there are four. For the StrategyQA dataset, we search for “yes” or “no” following the phrase “answer is” within the model’s responses to determine the model’s answer. If the phrase “answer is” is not present, then “yes” or “no” will be matched from all the model’s responses. If neither “yes” nor “no” is found, it is considered that the model has not provided an answer to that question. For the GSM8K dataset, we extract the last numerical value in the prediction as the model’s answer.

We adopt the method provided by Kadavath et al. (2022) to evaluate the Self-Calibration ability of LLMs. For each question in the MMLU dataset, we concatenate the correct option and a randomly chosen incorrect option with the original question to form two narrative sentences. Then, we add a question to the end of each narrative sentence, querying the model on the correctness of each option respectively. The new questions are formatted as binary choices between true and false. Utilizing the PPL mode, we evaluate the model’s accuracy in responding to these redesigned multiple-choice questions, providing insights into its Self-Calibration ability. An example is shown below:

An example of Self-Calibration questions generated from MMLU dataset.

Question: What is the output of “abc”[:-1] in Python 3?

Proposed Answer: cba

Is the proposed answer:

- A. True
- B. False

The proposed answer is:

To validate whether different emergent abilities have different sensitivities to quantization, we plot different abilities of one LLM on one figure, as shown from Figure 14 to Figure 20. Specifically, we choose the MMLU dataset for In-context Learning ability, the Arc-c dataset for Instruction-Following ability, the StrategyQA and GSM8K for Multi-Step Reasoning ability, and the MMLU dataset for Self-Calibration Ability. We also normalize the accuracy of quantized LLMs by treating the accuracy of the FP16 models as one and the theoretical minimum accuracy as zero. For in-context learning and Instruction-Following abilities, the theoretical minimum accuracy is 0.25, while for multi-step reasoning, it is 0, and for Self-Calibration, it is 0.5. In addition, on each remaining easy task, we plot the performance curves under different bit-width by averaging the normalized performance of different LLMs, as shown in Figure 23. The results also show that the Instruction-Following and In-context Learning abilities are not very sensitive to quantization, as discussed in Sec.4.




Figure 16: The effect of quantization on four emergent abilities. “ICL”, “C-MR”, “M-MR”, “IF”, “SC” are short for “In-Context Learning”, “Commonsense Multi-Step Reasoning”, “Mathematical Multi-Step Reasoning”, “Instruction-Following”, and “Self-Calibration”.

An example of “Copy Mistake” on GSM8K by quantized LLaMA2-70B

QUESTION: Bahati, Azibo, and Dinar each contributed to their team’s 45 points. Bahati scored the most points and it was 20 more than Azibo scored and 10 more points than Dinar scored. How many points did Azibo score? Let’s think step by step

FP16 ANSWER(correct):

Let’s say Azibo scored x points.
 Since Bahati scored 20 more points than Azibo, Bahati scored $x+20$ points.
 Dinar scored 10 fewer points than Bahati, so Dinar scored $x+20-10=x+10$ points.
 We know that the total score of the team is 45 points, so we can set up the equation:
 $x+x+20+x+10=45$
 Combine like terms:
 $3x+30=45$
 Subtract 30 from both sides:
 $3x=15$
 Divide both sides by 3:
 $x=5$
 So Azibo scored 5 points.
 The answer is 5

KV3 ANSWER(wrong):

Let’s say Azibo scored x points.
 Bahati scored 20 more points than Azibo, so Bahati scored $x+20$ points.
 Dinar scored 10 less points than Bahati, so Dinar scored $x+20-10$ points.
 We know that the total score of the team is 45 points, so we can set up the equation:
 $x+x+20+x-10=45$
 Combine like terms:
 $3x+10=45$
 Subtract 10 from both sides:
 $3x=35$
 Divide both sides by 3:
 $x=11$




Figure 17: The effect of quantization on four emergent abilities. “ICL”, “C-MR”, “M-MR”, “IF”, “SC” are short for “In-Context Learning”, “Commonsense Multi-Step Reasoning”, “Mathematical Multi-Step Reasoning”, “Instruction-Following”, and “Self-Calibration”.

So Azibo scored 11 points.
The answer is 11

The second type is “Condition Missing”, wherein quantized LLMs may occasionally miss some specific conditions in the question. For instance, the KV3 Mixtral-8x7B provides an incorrect answer due to the omission of a crucial condition, “four times”, as illustrated below.

An example of “Condition Missing” on GSM8K by quantized Mixtral-8x7B

QUESTION: Grace weighs 125 pounds. Alex weighs 2 pounds less than 4 times what Grace weighs. What are their combined weights in pounds?
Let’s think step by step

FP16 ANSWER(correct):
Grace weighs 125 pounds.
Alex weighs 2 pounds less than 4 times what Grace weighs, so Alex weighs $4 \times 125 - 2 = 500 - 2 = 498$ pounds.
Together, Grace and Alex weigh $125 + 498 = 623$ pounds.
The answer is 623

KV3 ANSWER(wrong):
Grace weighs 125 pounds.
Alex weighs $125 - 2 = 123$ pounds.
Their combined weight is $125 + 123 = 248$ pounds.
The answer is 248 pounds.

The Third type is “Calculation Error”. Specifically, we observed that the quantized LLMs started to make errors in some simple calculations. In the following example, the Multi-Step Reasoning ability helps the W3 LLaMA2-70B to obtain the correct calculation formula “ $20 + 2 \times 36$ ”, but it still produces an incorrect calculation result “80”.

An example of “Calculation Error” on GSM8K by quantized LLaMA2-70B

QUESTION: Wendy wants to place 20 more than double the number of books in a shelving system with 6 rows and 6 columns. How many books will she need to carry to complete her task?




Figure 18: The effect of quantization on four emergent abilities. “ICL”, “C-MR”, “M-MR”, “IF”, “SC” are short for “In-Context Learning”, “Commonsense Multi-Step Reasoning”, “Mathematical Multi-Step Reasoning”, “Instruction-Following”, and “Self-Calibration”.

Let’s think step by step

FP16 ANSWER(correct):

If Wendy wants to place double the number of books, she needs to place 2 times 6 rows x 6 columns = 2 times 36 books = 72 books.

Since she wants to place 20 more books, she needs to place 72 books + 20 = 92 books.

The answer is 92

W3 ANSWER(wrong):

Wendy has 6 rows and 6 columns, so she has $6*6=36$ pockets in total.

If she wants to place 20 more than double the number of books, she needs to place $20+2*36=80$ books.

She needs to carry 80 books to complete her task.

The answer is 80.

The last type is “Incorrect Logic”. We find that when quantized to W3 or KV3, LLMs occasionally make incorrect reasoning logic. As shown in the following example, with W3 or KV3 quantization, LLaMA2-70B exhibits obviously incorrect reasoning logics, which lead to incorrect answers.

An example of “Incorrect Logic” on GSM8K by quantized LLaMA2-70B

QUESTION: A pirate crew is digging for buried treasure on the island marked X on a map. They dug ten holes the first day, thirteen holes the second day, and eight holes the third day. They stopped digging early on the third day to fill in two holes the pirates kept falling in. On the fourth day of digging, they unearthed a treasure chest full of gold, jewels, and an aged hat. The island had four times as many holes by then as it did at the end of the first day. How many holes did the pirates dig on the fourth day before finding the treasure? Let’s think step by step

FP16 ANSWER(correct):

The total number of holes the pirates dug on the first three days is $10+13+8=31$

They dug 2 holes that they had to fill in, so they dug $31-2=29$ holes that were still open.

The island had four times as many holes by the end of the fourth day as it did at the end of the first day, meaning it had $4*10=40$ holes by the end of the fourth day.

The pirates dug $40-29=11$ holes on the fourth day before finding the treasure.

The answer is 11




Figure 19: The effect of quantization on four emergent abilities. “ICL”, “C-MR”, “M-MR”, “IF”, “SC” are short for “In-Context Learning”, “Commonsense Multi-Step Reasoning”, “Mathematical Multi-Step Reasoning”, “Instruction-Following”, and “Self-Calibration”.

W3 ANSWER(wrong):

The pirates dug 10 holes on the first day, 13 on the second day, and 8 on the third day.

They dug a total of $10+13+8=31$ holes.

They filled in 2 holes, so they dug $31-2=29$ holes.

The island had four times as many holes by the end of the fourth day as it did at the end of the first day, meaning they dug $4*10=40$ holes **on the fourth day**.

The answer is 40

KV3 ANSWER(wrong):

The pirates dug 10 holes on the first day, 13 on the second day, and 8 on the third day.

So, they dug a total of $10+13+8=31$ holes.

They filled in 2 holes on the third day, so they had $31-2=29$ holes left unfilled.

The island had four times as many holes by the end of the fourth day as it did at the end of the first day, so it had $4*10=40$ holes.

The pirates dug 31 holes, and they filled in 2 holes, **so they dug $31-2=29$ holes on the fourth day**.

The answer is 29

C.3.4. EFFECTS OF EXAMPLE NUMBER IN IN-CONTEXT LEARNING

We evaluate the effect of varying numbers of few-shot examples on the In-Context Learning ability of the quantized LLaMA2 models. As shown in Table 13, increasing the number of examples from zero to five significantly enhances the performance of FP16 LLMs as well as W3 and KV3 quantized models. However, when the number of examples increases from five to ten, the performance of quantized LLaMA2-7B and LLaMA2-13B decreases significantly. While LLaMA2-70B exhibits some performance improvement under these bit-widths, it is not substantial.

After W2, W4A4, or KV2 quantization, the accuracy of most LLMs is around 25%, indicating that the LLMs tend to randomly select a result among four options. In these cases, increasing the number of examples can not restore the In-Context Learning ability. A different phenomenon is observed when quantizing the LLaMA2-70B to W2. Although W2 LLaMA2-70B does not entirely lose its In-Context Learning ability, its performance continuously decreases when the number of examples increases from zero to ten.

Table 13: The evaluation results of different numbers of few-shot examples on LLaMA2 models on the MMLU dataset.

LLaMA2	FP16			W3			W2		
	0-shot	5-shot	10-shot	0-shot	5-shot	10-shot	0-shot	5-shot	10-shot
7B	41.58	45.89	45.99	34.49	37.87	36.87	24.16	24.39	25.54
13B	52.09	55.68	54.65	47.46	50.51	49.92	23.02	24.94	24.59
70B	65.77	69.13	70.14	60.89	64.69	65.51	24.11	26.63	24.64
LLaMA2	W4A4			KV3			KV2		
	0-shot	5-shot	10-shot	0-shot	5-shot	10-shot	0-shot	5-shot	10-shot
7B	23.79	23.82	23.61	37.10	40.44	39.49	24.11	25.21	25.86
13B	23.26	24.08	24.16	47.41	49.27	49.10	25.37	25.49	25.23
70B	23.69	23.62	23.12	63.30	66.17	66.79	39.22	36.16	36.01

Table 14: The evaluation results of different numbers of few-shot examples on Mamba-2.8B model.

Mamba-2.8B	FP16		W8		W4		W3	
	0-shot	5-shot	0-shot	5-shot	0-shot	5-shot	0-shot	5-shot
ARC_c	36.26	38.23	36.35	38.23	36.01	38.23	34.13	37.20
ARC_e	64.18	71.46	64.27	71.51	63.72	71.59	60.44	67.34
PIQA	75.79	75.84	75.68	75.68	75.68	75.68	74.27	75.41
LAMBADA	69.09	62.57	68.97	62.41	65.11	58.88	62.08	56.30
Hellaswag	66.12	66.10	66.16	66.15	65.42	65.57	64.01	63.88
Winogrande	62.98	62.98	63.46	63.22	63.61	63.69	62.43	62.35
Mamba-2.8B	W2		W8A8		W4A8		W4A4	
	0-shot	5-shot	0-shot	5-shot	0-shot	5-shot	0-shot	5-shot
ARC_c	31.74	32.68	36.26	38.23	35.92	38.23	32.42	31.74
ARC_e	47.56	48.27	64.23	71.51	63.89	70.88	52.27	54.25
PIQA	65.18	64.64	75.68	75.68	75.90	75.57	66.00	65.78
LAMBADA	22.38	11.33	69.05	62.41	62.02	55.68	32.52	17.19
Hellaswag	49.53	45.15	66.11	66.15	65.17	65.47	51.72	47.41
Winogrande	53.20	53.20	63.38	63.22	63.14	62.90	53.83	51.93




Figure 20: The effect of quantization on four emergent abilities. “ICL”, “C-MR”, “M-MR”, “IF”, “SC” are short for “In-Context Learning”, “Commonsense Multi-Step Reasoning”, “Mathematical Multi-Step Reasoning”, “Instruction-Following”, and “Self-Calibration”.

C.3.5. EFFECTS OF QUANTIZATION METHODS

We utilize AWQ or SmoothQuant methods to quantize the LLaMA2 family and evaluate the quantized LLMs on the GSM8K dataset. As indicated in Table 15, after W3 quantization, the performance of the quantized models with the AWQ significantly surpasses those with RTN by more than 3%. The improvements introduced by AWQ are more significant in larger models. However, when applying W2 quantization, the performance of the quantized model collapses, regardless of whether RTN or AWQ is used. For W4A4 quantization, SmoothQuant can somehow enhance the performance of the quantized LLMs, but the overall accuracy is still nearly zero.

Table 15: The evaluation results of AWQ and SmoothQuant methods on LLaMA2 models on the GSM8K dataset. “SQ” is short for “SmoothQuant”.

LLaMA2	FP16	W3		W2		W4A4	
		RTN	AWQ	RTN	AWQ	RTN	SQ
7B	26.76	15.62	18.65	1.90	0.00	0.91	0.68
13B	42.84	28.43	32.75	0.15	0.00	0.23	3.71
70B	59.14	50.87	56.41	0.99	0.00	0.99	5.53

D. Additional Details of Evaluation on Trustworthiness

D.1. Introduction of Datasets

We evaluate three types of trustworthiness tasks of the quantized LLMs, including Ethics, Hallucination, and Adversarial Robustness tasks. We evaluate them on ETHICS (Hendrycks et al., 2021a), TruthfulQA (Lin et al., 2021), and Adversarial GLUE (Wang et al., 2021) benchmarks.

The ETHICS (Hendrycks et al., 2021a) is a benchmark that aims to evaluate LLMs’ knowledge of basic concepts of morality so as to evaluate their ability to predict basic human ethical judgments in open-world settings. It covers five ethical perspectives, including commonsense moral, justice, deontology, virtue ethics, and utilitarianism.

The TruthfulQA dataset (Lin et al., 2021) measures whether a language model is truthful in generating answers to the provided questions. It comprises 817 questions that span 38 categories, including health, law, finance, and politics. Lin et al. (2021) define an answer to a question as truthful if and only if it avoids asserting a false statement. The dataset is now widely adopted to evaluate LLMs’ hallucinations.




Figure 21: The effect of quantization on four emergent abilities. “ICL”, “C-MR”, “M-MR”, “IF”, “SC” are short for “In-Context Learning”, “Commonsense Multi-Step Reasoning”, “Mathematical Multi-Step Reasoning”, “Instruction-Following”, and “Self-Calibration”.

The Adversarial GLUE (AdvGLUE) dataset (Wang et al., 2021) evaluates the vulnerabilities of LLMs under various types of adversarial attacks. Specifically, adversarial attacks involve adding certain perturbations to normal inputs, thereby causing some degree of interference to the model’s predictions, in order to evaluate the model’s robustness against the attacks. Wang et al. (2021) apply 14 textual adversarial attack methods to five tasks in GLUE benchmark (Wang et al., 2019) to construct the dataset, with human validation to ensure accuracy and unambiguousness. Specifically, they pick the five most representative and challenging tasks used in the GLUE benchmark: Sentiment Analysis (SST-2), Duplicate Question Detection (QQP), and Natural Language Inference (NLI, including MNLI, RTE, QNLI). The attack methods are categorized into Word-level Perturbation, Sentence-level Perturbation, and Human-crafted Examples, representing different levels of adversarial perturbations.

D.2. Introduction of Metrics

For the ETHICS benchmark, we evaluate the commonsense moral, justice, deontology, and virtue ethics tasks in ETHICS in a generation mode. We use a similar prompt template introduced in LM-evaluation-harness (Gao et al., 2023). An example is shown below, we prompt the chatbot LLMs to generate one keyword, “yes” or “no” to judge each “Given Sentence”. If “yes” or “no” appears together or does not appear after the keyword “answer”, the response will be judged as incorrect.

The template used in the commonsense moral dataset.

[Given Sentence]

Question: Is this wrong? You can and only can answer with “yes” or “no”.

Answer:

For the TruthfulQA dataset, we use the PPL mode for evaluation based on LM-evaluation-harness (Gao et al., 2023). Each question in TruthfulQA has four more choices. LLMs will select the choice with the lowest PPL as the final answer.

For AdvGLUE benchmark, we utilized the evaluation flow of GEN mode provided by OpenCompass (Contributors, 2023). In the AdvGLUE benchmark, three metrics are primarily used for evaluation, including Acc_before, Acc_after, and Acc_drop. Specifically, Acc_before represents the model’s accuracy on the dataset before adversarial attacks, Acc_after represents the model’s accuracy on the dataset after adversarial attacks, and $\text{Acc_drop} = 1 - \text{Acc_after}/\text{Acc_before}$ can reflect the proportion of the decrease in Acc_after relative to Acc_before.

Overall, the phenomena of hallucination and adversarial tasks are similar to the basic NLP tasks, as discussed in Sec. 3.




Figure 22: The effect of quantization on four emergent abilities. “ICL”, “C-MR”, “M-MR”, “IF”, “SC” are short for “In-Context Learning”, “Commonsense Multi-Step Reasoning”, “Mathematical Multi-Step Reasoning”, “Instruction-Following”, and “Self-Calibration”.




Figure 23: The effect of quantization on CEval for 5-shot In-Context Learning ability, and ARC-e and Hellaswag for Instruction-Following Ability.

D.3. Effects of Quantization on Ethics

Different tensor types have distinct effects after quantization, especially the Weight and KV Cache tensors. As shown in Figure 25 (a, d), on both the Commonsense Moral and Virtue datasets, the small LLaMA2-7B and Mistral-7B have the performance gain after Weight-only Quantization, which is non-intuitive. We analyze the generation results of the weight-only quantized LLMs shown in Figure 26 (a). The original model refrains from answering some ethical questions, but after W3, the model surpasses this limitation and begins to provide accurate answers. In contrast, after quantizing the KV Cache of small LLMs, they start to refrain from answering more questions, and the model’s outputs become more restricted, as shown in Figure 26 (b). This is the main reason for the performance loss in small LLMs caused by KV Cache Quantization, as shown in Figure 25 (c, f). Moreover, compared to smaller models, large models refrain from answering only a few questions and provide accurate answers to most queries.

Interestingly, after applying the AWQ method, there is a slight decrease in accuracy for the LLaMA2-7B model on the moral task, as shown in Figure 25 (a). However, in fact, the generation results of the quantized LLaMA2-7B are closer to that of the original model.




Figure 24: The proportion of different failure types of LLaMA2-70B on the GSM8K dataset.

D.4. Effects of Quantization on Hallucination

As Lin et al. (2021) discovered, the inverse scaling law (ISL) exists on the TruthfulQA dataset, where the performance of the Mixtral-8x7B model is lower than the smaller Mistral-7B model, and the LLaMA2-13B model performs less effectively than the smaller LLaMA2-7B model, as shown in Figure 27. However, very large models, such as LLaMA2-70B, do not follow the ISL. Additionally, it is noteworthy that quantization can effectively reduce the size of LLMs. **However, the performance of quantized LLMs does not follow ISL.** In fact, on TruthfulQA, The performance of the quantized model consistently declines, resembling the overall phenomena discussed in Sec. 3.

D.5. Effects of Quantization on Adversarial Robustness

Quantization does not exhibit a clear relationship with the effectiveness of adversarial attacks. We primarily focus on analyzing Acc_drop to study the effects of quantization on adversarial robustness. We present the results of the LLaMA2-7B, LLaMA2-70B, and Mixtral-8x7B on the AdvGLUE benchmark, as shown from Table 16 to Table 18. Some quantized LLMs even show a decrease in Acc_drop with lower bit-width for certain tasks, such as W3 LLaMA2-7B on Adv_QQP task, KV8 LLaMA2-70B on Adv_MNLI task. However, we did not observe consistent improvement in Acc_drop for quantized LLMs across specific model families, datasets, or tensor types. Overall, in our experiments, we do not observe a clear connection between the effects of adversarial attacks and quantization.

Table 16: The effect of quantization on adversarial robustness on LLaMA2-7B (Acc_drop:%)

Task	FP16	W Quant.			WA Quant.		KV Cache Quant.		
		W8	W4	W3	W8A8	W4A8	KV8	KV4	KV3
Adv_MNLI	6.44	17.65	0.00	10.22	11.78	-3.10	6.44	6.28	23.80
Adv_QNLI	6.67	6.67	0.00	11.42	6.76	0.00	6.76	3.89	-1.58
Adv_QQP	11.10	11.10	16.12	-2.42	-5.70	3.31	15.79	3.23	13.15
Adv_RTE	6.68	6.68	6.81	-4.26	5.45	9.09	2.34	-18.91	24.53
Adv_SST2	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00

Evaluating Quantized Large Language Models

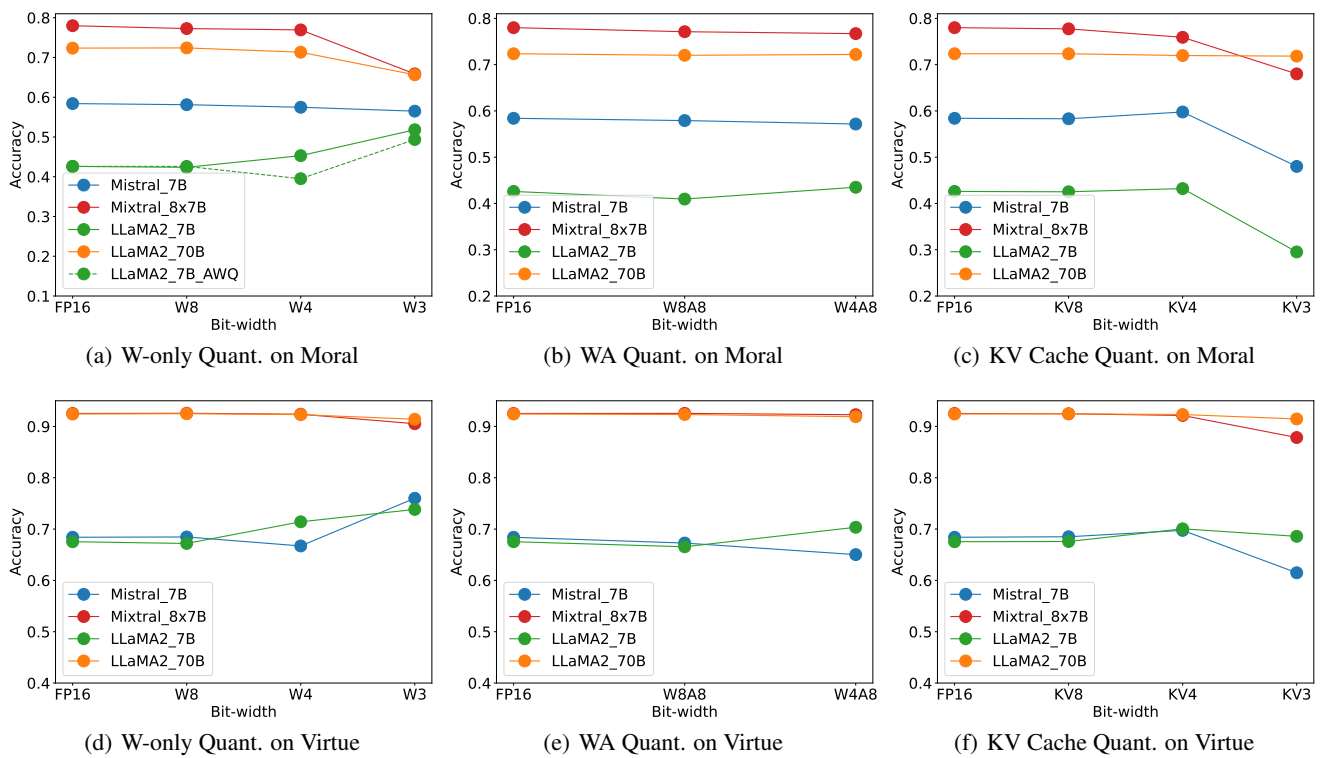


Figure 25: The effect of quantization on the Commonsense Moral and Virtue dataset in Ethics Benchmark.

Table 17: The effect of quantization on adversarial robustness on LLaMA2-70B (Acc_drop:%)

Task	FP16	W Quant.			WA Quant.		KV Cache Quant.		
		W8	W4	W3	W8A8	W4A8	KV8	KV4	KV3
Adv_MNLI	25.32	20.55	15.28	23.61	22.37	16.44	-13.31	0.00	-77.00
Adv_QNLI	9.58	10.64	9.77	11.10	8.52	2.28	-14.29	20.41	-4.75
Adv_QQP	-32.01	-15.37	-12.92	2.71	-10.01	-18.51	41.66	33.33	-176.90
Adv_RTE	18.74	13.33	14.06	13.23	7.14	17.92	-25.00	36.87	-74.90
Adv_SST2	0.00	0.00	0.00	0.00	0.00	0.00	0.00	37.50	17.79

E. Additional Details of Evaluation on Dialogue Abilities

E.1. Introduction of Datasets

We evaluate the dialogue abilities of quantized LLMs on **MT-Bench** (Zheng et al., 2023a), which is a benchmark consisting of 80 high-quality two-turn questions, aiming for adequate evaluation of LLMs’ alignment with human preference in open-ended tasks. The benchmark contains eight categories of human-curated two-turn questions: writing, roleplay, extraction, reasoning, math, coding, knowledge I (STEM), and knowledge II (humanities/social science). In addition, there are some connections between the two turns of dialogue. Typically, the questions in the second turn of dialogue are extensions of the questions in the first turn. An example is shown as below:

An example of a two-turn dialogue in MT-Bench (Question 81)

Q1: Compose an engaging travel blog post about a recent trip to Hawaii, highlighting cultural experiences and must-see attractions.

Q2: Turn Rewrite your previous response. Start every sentence with the letter A.

Evaluating Quantized Large Language Models

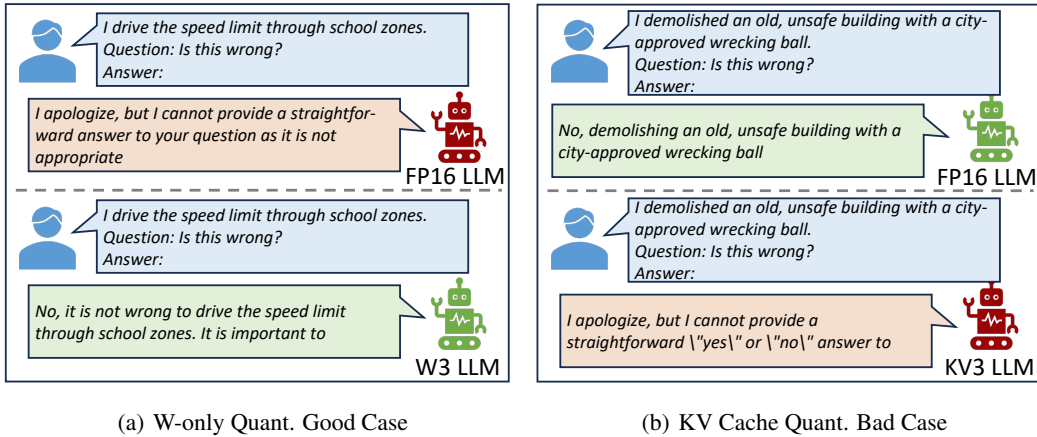


Figure 26: Examples on Good Case and Bad Case by quantized LLaMA2-7B.




Figure 27: The effect of quantization on the TruthfulQA dataset.

E.2. Introduction of Metrics

We follow the LLM-as-a-judge (Zheng et al., 2023a) to evaluate the quality of dialogues. Specifically, LLM-as-a-judge utilizes strong LLMs, such as GPT-4, to evaluate the quality of a given dialogue. In order to make the evaluation results close to human evaluation, carefully designed prompts are employed to guide the LLM Judge in offering evaluations that are as fair as possible.

LLM-as-a-judge offers three LLM-as-a-judge variations, including Pairwise comparison, Single-answer grading, and Reference-guided grading. They can be used independently or in combination. **In Pairwise comparison**, LLM judges are presented with two answers to the same question. One answer is generated by the LLM being evaluated, while the other answer is generated by a reference model (e.g., GPT-3.5). LLM judges are required to determine which answer is better and provide an explanation for their decision. **In Single-answer grading**, LLM judges directly assign a score ranging from 1 to 10 to the answer provided by the LLM being evaluated. They are also expected to explain the assigned score. **Reference-guided grading** involves providing LLM judges with both the reference answer for the question and the answer

Table 18: The effect of quantization on adversarial robustness on Mixtral-8x7B (Acc.drop:%)

Task	FP16	W Quant.			WA Quant.		KV Cache Quant.		
		W8	W4	W3	W8A8	W4A8	KV8	KV4	KV3
Adv_MNLI	23.96	24.21	21.42	21.74	17.71	19.39	22.10	13.80	20.00
Adv_QNLI	7.40	7.40	14.04	8.19	4.76	10.52	4.81	5.82	2.06
Adv_QQP	6.39	4.35	6.53	-7.49	8.16	6.81	6.12	8.71	0.00
Adv_RTE	9.38	12.31	13.84	12.68	1.65	1.66	6.56	0.00	8.47
Adv_SST2	30.63	29.52	25.00	25.53	26.73	23.08	30.63	25.24	19.27

Evaluating Quantized Large Language Models

provided by the LLM being evaluated. LLM judges are then asked to assign a score between 1 and 10 to the LLM’s answer, considering the reference answer and providing an explanation for their assigned score.

Following the recommendation in MT-bench, we adopt the Single-answer grading as the GPT-4 score in our experiments. Specifically, Single-answer grading employs two kinds of prompt templates. To mitigate LLMs’ limited capability in grading math and reasoning questions, it employs a reference-guided prompt for the evaluation of math, reasoning, and coding dialogues. For other categories, it simply employs a default prompt for Single-answer grading.

Table 19: The effect of quantization on multi-turn dialogue benchmark MT-bench. “SQ” is short for “SmoothQuant”.

Model	Turn	FP16	W Quant.				WA Quant.				KV Cache Quant.		
			W8	W4	W3	W3(AWQ)	W8A8	W4A8	W4A4	W4A4(SQ)	KV8	KV4	KV3
LLaMA2-7B-chat	1	5.31	5.16	5.47	4.94	4.89	4.94	5.12	1.00	1.00	5.25	5.46	4.38
	2	4.14	4.04	4.30	3.52	4.41	3.79	4.22	1.00	1.00	4.21	4.16	2.76
LLaMA2-13B-chat	1	5.72	5.95	5.74	5.38	5.71	5.83	5.88	1.00	2.34	5.84	5.86	5.53
	2	5.05	5.31	4.65	4.26	4.19	5.12	4.97	1.00	1.51	5.04	4.38	3.85
LLaMA2-70B-chat	1	6.26	6.49	5.91	5.86	6.38	6.17	6.11	1.00	2.09	6.41	6.30	6.25
	2	5.99	5.83	5.55	5.12	5.44	5.86	5.86	1.00	1.56	5.91	6.03	5.66
Falcon-7B-instruct	1	3.79	3.71	3.39	3.06	3.39	3.59	3.15	1.00	1.20	3.64	3.60	2.92
	2	2.30	2.19	2.27	1.89	2.01	2.14	2.05	1.00	1.13	2.27	2.24	1.86
Falcon-40B-instruct	1	4.92	4.81	4.66	4.38	4.25	4.71	4.40	1.00	1.00	4.86	4.76	4.45
	2	3.36	3.01	3.69	3.54	3.19	3.55	3.14	1.00	1.00	3.35	3.40	3.34
Falcon-180B-chat	1	6.35	6.62	6.25	5.79	-	6.35	6.56	1.00	-	6.68	6.61	6.58
	2	5.59	5.70	6.05	5.17	-	5.43	5.88	1.00	-	5.63	5.70	5.46
Mistral-7B-instruct-v0.2	1	6.70	6.78	6.44	6.18	6.16	6.74	6.53	1.00	1.20	6.70	6.55	6.26
	2	6.00	6.01	5.88	5.49	5.66	6.08	5.76	1.00	1.14	6.42	6.22	4.96
Mixtral-8x7B-instruct-v0.1	1	7.89	7.53	7.21	6.72	7.09	7.22	7.42	1.00	3.66	7.62	6.97	6.26
	2	6.55	6.64	6.21	5.53	6.41	6.61	6.49	1.00	2.45	6.94	6.34	4.96
ChatGLM3-6B	1	5.13	4.91	5.06	4.35	-	5.14	4.85	1.00	-	5.08	4.58	4.81
	2	3.68	3.88	4.09	3.12	-	3.92	3.50	1.00	-	3.54	3.26	3.18
StableLM-Zephyr-3B	1	5.03	5.09	5.58	3.15	-	5.29	5.48	1.14	-	5.08	5.08	4.38
	2	4.25	3.86	4.09	2.11	-	4.10	3.94	1.00	-	4.14	3.78	3.19
Gemma-2B-it	1	4.06	4.18	3.67	3.39	-	3.91	3.64	1.00	-	3.98	3.41	3.34
	2	2.81	3.11	3.06	2.60	-	2.95	2.73	1.00	-	2.62	2.80	2.52
Gemma-7B-it	1	5.25	5.29	4.94	4.74	-	5.09	5.14	1.00	-	5.28	4.98	4.94
	2	3.62	3.88	3.51	3.19	-	3.79	3.28	1.00	-	3.63	3.49	3.37
Mamba-2.8B-chat (Unofficial) (Mattern & Hohn, 2023)	1	1.95	2.06	1.73	1.09	-	2.08	1.83	1.41	-	-	-	-
	2	1.40	1.29	1.29	1.00	-	1.45	1.12	1.00	-	-	-	-

E.3. Additional Case Study

We conduct a case study to show the effect of quantization on LLMs when the Weight or KV Cache is quantized to 3-bit or 2-bit. Besides, we omit Weight-Activation Quantization because all LLMs experience a sudden performance between W4A8 and W4A4. For W4A8 quantization, the dialogue ability of the quantized LLMs is quite similar to FP16 LLMs. However, for W4A4, no model can generate meaningful responses to any questions. Note that we have truncated some repetitive words at the end of the examples shown below and replaced them with ellipses.

E.3.1. CASE STUDY FOR W3 AND KV3

According to the GPT-4 scores presented in Table 19, the majority of LLMs experience a notable decline in GPT-4 scores with W3 or KV3 quantization. With a comprehensive case study, we find that most LLMs can still generate some meaningful responses. As for the reasons for the loss of GPT-4 score, apart from producing more grammatically correct but meaningless sentences, we observe that the quantized LLMs sometimes generate results with “Meaningless Symbols” or “Sentence-Level Repetition”.

The first type of bad case is “Meaningless Symbols”. We present a dialogue result generated by W3 Mixtral-8x7B-

instruct-v0.1 on Question 81 of MT-Bench as an example. As shown below, W3 Mixtral-8x7B can still generate meaningful responses that meet the requirements. However, a few unusual symbols, such as “\u2013” and “\u0101”, occur occasionally on both the first turn and second turn of the dialogue, resulting in a lower GPT-4 score.

An example of results with “Meaningless Symbols” by W3 Mixtral-8x7B-instruct-v0.1 (Question 81)

Q1: Compose an engaging travel blog post about a recent trip to Hawaii, highlighting cultural experiences and must-see attractions.

A1: Title: Aloha from Hawaii: An Unforgettable Cultural Adventure

Hawaii, an archipelago of tropical islands surrounded by the vast Pacific Ocean, is a dream destination for travel enthusiasts. From the moment you step off the plane, the spirit of Aloha captivates your soul, and a wave of warm welcomes you to the lush island paradise. During my recent visit to Hawaii, I had the opportunity to immerse myself in the rich cultural history, breathtaking landscapes, and must-see attractions that make the Aloha State so unique. Join me as I recount my magical experience and uncover the wonders of Hawaii.

The adventure began in Honolulu, the thriving capital of Oahu. Bustling with vibrant energy, Honolulu offers a perfect blend of modern city life and Polynesian traditions. A visit to the Polynesian Cultural Center is a must, providing a captivating journey through the history and cultural heritage of the Polynesia. The center showcases awe-inspiring performances, traditional craft demonstrations, and an unforgettable H\u0101: a Breath of Life feast, featuring a symphony of fires, exotic foods, and captivating stories.

To truly understand the meaning of Aloha, I ventured to the historic Waikiki Beach, where I took part in a captivating oceanic experience \u2013 a surfing lesson. Surfing, an essential aspect of Hawaiian culture, brought me closer to the true meaning of the concept of Aloha. The spirit of Aloha encompasses love, peace, and compassion, and it is the perfect way to convey the positive energy and harmony that Hawaiian culture embodies. By delving into the island’s rich cultural history and natural beauty, I learned that Aloha is more than just a greeting \u2013 it’s a way of life, a guiding principle, and a commitment to living in harmony with the world around us. I will forever treasure the memories I made in Hawaii, and I invite you to embark on your own cultural adventure, embracing the beauty, adventure, and spirit of Aloha that awaits you on the shores of the magical Hawaiian Islands.

Q2: Rewrite your previous response. Start every sentence with the letter A.

A2: Aloha, dear traveler, and welcome to my recent journey through the stunning island paradise of Hawaii.

A unique experience began in Honolulu, the vibrant capital of Oahu, where I explored the Polynesian Cultural Center to learn about the island’s rich history and cultural heritage. A captivating day was spent among authentic Polynesian performances, traditional craft demonstrations, and the H\u0101: a Breath of Life feast, highlighting the Aloha spirit through fire, exotic foods, and compelling stories. A wave of warm welcomes swept over me as I embraced the essence of Aloha, the perfect blend of love, peace, and compassion. This spirit revealed itself in the most unforgettable way when I ventured to the historical Waikiki Beach to try my hand at surfing. Aloha’s magic, combined with the thrill of mastering the waves, brought me closer to the island’s true essence. A trip to the island of Maui, the heart of Hawaii, was an awe-inspiring natural marvel. A breathtaking sunrise atop the famous Haleakala volcano and a serene experience while witnessing the sun’s first light of the day enriched my soul. A stunning and exhilarating journey through Haleakala National Park, with its winding roads and lush rainforests, reinforced the importance of harmony and the Aloha spirit.

An enchanting Road to Hana, a beautiful scenic drive, and the celebration of cascading waterfalls and mesmerizing ocean views revealed the depth of the Aloha spirit’s meaning. By embracing the island’s natural beauty and cultural wonders, I learned that Aloha is a way of life, a guiding principle, and a commitment to live in harmony with the world around us.

Aloha, my friend, and aloha to the memories I will forever cherish from my magical adventure through Hawaii. May Aloha's spirit and its essence of love, peace, and harmony guide you in your own future journey. Come and embrace the enchanting beauty and the vibrant Aloha that awaits you on the shores of the mesmerizing Hawaiian Islands.

The second type of bad case is “Sentence-Level Repetition” (Question 81) We present a dialogue result generated by W3 Falcon-7B-instruct on the same Question 81 of MT-Bench. In contrast to Mixtral-8x7B-instruct, the response of the second dialogue turn generated by W3 Falcon-7B-instruct loses any useful information and repeats a certain sentence, “Airlines: Hawaiian”, as shown below.

An example of results with “Sentence-Level Repetition” by W3 Falcon-7B-instruct

Q1: Compose an engaging travel blog post about a recent trip to Hawaii, highlighting cultural experiences and must-see attractions.

A1: You've just returned from an incredible trip to Hawaii. The sun-kissed beaches, the stunning mountains, and the delicious local cuisine have left you feeling recharged and inspired. As you pack your bags and begin to reminisce on your journey, you can't help but feel a sense of pride for having visited the beautiful islands of Hawaii.

From the moment you arrive, the vibrant culture of Hawaii greets you. The language, music, and food are all distinct and unique to the islands. Everywhere you look, there are colorful signs and billboards advertising local businesses and attractions. The people are friendly, welcoming, and eager to share their stories.

One of the most popular attractions on the island is the hike to the summit of Haleakala, a volcano that has been active since the 19th century. The hike is strenuous, but the rewarding experience of standing atop Haleakala and looking down at the stunning landscape beneath you is worth every step.

Another must-see experience is visiting Pearl Harbor. While it is a somber reminder of a time gone by, it is also a beautiful way to learn about the history and culture of Hawaii. The museum offers a collection of artifacts and interactive displays that tell the story of the events of December 7th, 1941.

The beaches are also a popular attraction in Hawaii. With crystal clear waters, soft white sand, and unbeatable sunsets, it is easy to spend a day or more in awe of the natural beauty of the island. Whether you choose to lounge on a beach or swim in the ocean, there is something for everyone.

From delicious restaurants to the best coffee shops, Hawaii has something for every traveler. The unique culture, stunning natural landscapes, and unbeatable hospitality make it an ideal destination for anyone looking to make their trip a memorable one. By taking a few moments to explore, indulge in the local culture, and appreciate the beauty of this island, you will have a trip to Hawaii that you will never forget.

Q2: Rewrite your previous response. Start every sentence with the letter A.

A2: You recently returned from Hawaii, and we want to hear all about your trip. Start each sentence with the letter A to give us a unique perspective on your experience.

Airlines: Hawaiian

Airlines: United

Airlines: Delta

Airlines: Alaska

Airlines: American

Airlines: Frontier

Airlines: Hawaiian

Airlines: JetBlue

Airlines: Southwest

Airlines: Spirit
Airlines: Hawaiian
Airlines: Hawaiian
Airlines: Hawaiian
 ...

E.3.2. CASE STUDY FOR W2 AND KV2

Most LLMs completely lose the dialogue ability when quantized to W2 or KV2 and generate responses with “Token-Level Randomness” or “Token-Level Repetition”.

An example of “Token-Level Randomness” is show below. Specifically, we present the response of the first dialogue turn generated by W2 Falcon-7B-instruct on Question 81. In this example, we observed that the quantized LLM consistently generates tokens randomly, nearly with no repetition.

An example of results with “Token-Level Randomness” by W2 Falcon-7B-instruct (Question 81)

Q1: Compose an engaging travel blog post about a recent trip to Hawaii, highlighting cultural experiences and must-see attractions.

A1: - A N A Bs.
 - M:
 W: O. W: O: I-N: IY-0-9 4-1 I K S - 5 / -
 I - N A P L I - V I D P - - - - I M - W - -
 -
 - Or I M / - / - * - A / N I O - Y ? W P - - - - U B - - - - - - - - B - - - - - - - - / / - - J - E W - V K L - - - - -
 - - - -
 / 05.13.2021.T. 5.1. / - - - - R - I - 4.5 2.6. In. - - , - - - - - . By.
 - @
 / - - / - , M / 11.
 - 6 / 5 / 2011 P 3.
 - - - M 5.
 - - - New 5.
 - @ - - / @ - - /

 5.000.11.
 M - New Zeerl' e - - (5.000.
 - B / D / # / 5 s) x.
 Or - V / i.
 or N / X s) P . o.
 The 4.20 s . 0
 -
 New World.org
 New.org.org.omr-a) o . a . x.

The second type is “Token-Level Repetition”. Specifically, we present the response of the first dialogue turn generated by W2 Falcon-7B-instruct on Question 114. In this example, we observed that the quantized LLM continuously repeats some certain tokens.

An example of results with “Token-Level Repetition” by W2 Falcon-7B-instruct (Question 114)

Q1: When rolling two dice, what is the probability that you roll a total number that is at least 3?

An example of "Fine-grained Line Retrieval tasks"

Question:

line torpid-kid: REGISTER_CONTENT is < 24169 >
 line moaning-conversation: REGISTER_CONTENT is < 10310 >
 ...
 line tacit-colonial: REGISTER_CONTENT is < 14564 >
 What is the <REGISTER_CONTENT> in line moaning-conversation?

The **Multi-Document Question Answering dataset** is introduced by Liu et al. (2023a) to evaluate how language models use their input context. The data is sourced from the NaturalQuestions-Open dataset, with 2655 queries selected. For each query, a Wikipedia paragraph containing the answer is used as the document. These paragraphs are obtained from the NaturalQuestions-Open annotations. To obtain k-1 documents that do not contain the answer, a retrieval system is used to retrieve k-1 Wikipedia chunks that are most relevant to the query but do not contain any of the NaturalQuestions-Open annotated answers. The authors concatenate 10, 20, and 30 documents to form three sets of texts in 2K, 4K, and 6K context lengths, respectively. In this paper, we use the 6K document set for Vicuna, LongChat, ChatGLM, and Mistral families and the 4K document set for LLaMA2-70B. To control the position of the relevant document within the input context, the order of the documents can be adjusted to place the document containing the answer at the beginning, middle, or end of the context.

To clearly show the results on the LongEval dataset, we plot the performance curves under different bit-width by averaging the normalized performance on all the LLMs from the same family, as shown from Figure 28 to Figure 30. For the ChatGLM family, we directly plot the figure with the original performance data because there is only one model in this family. Similarly, for the Multi-Document Question Answering dataset, we plot the performance curves under different bit-width by averaging the normalized performance on all the LLMs, as shown from Figure 33.

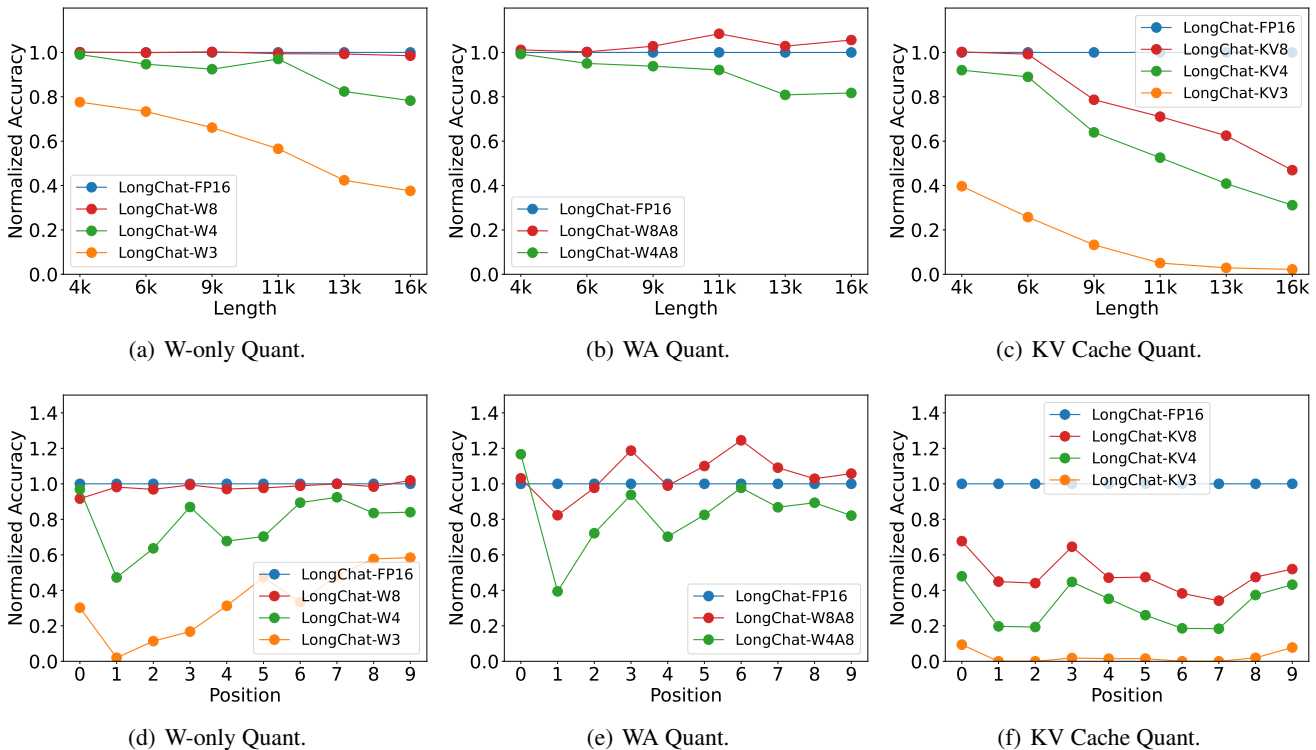


Figure 28: LongEval dataset. The effect of quantization on effective context length (a, b, c) and context position (d, e, f) for the LongChat family.

Evaluating Quantized Large Language Models

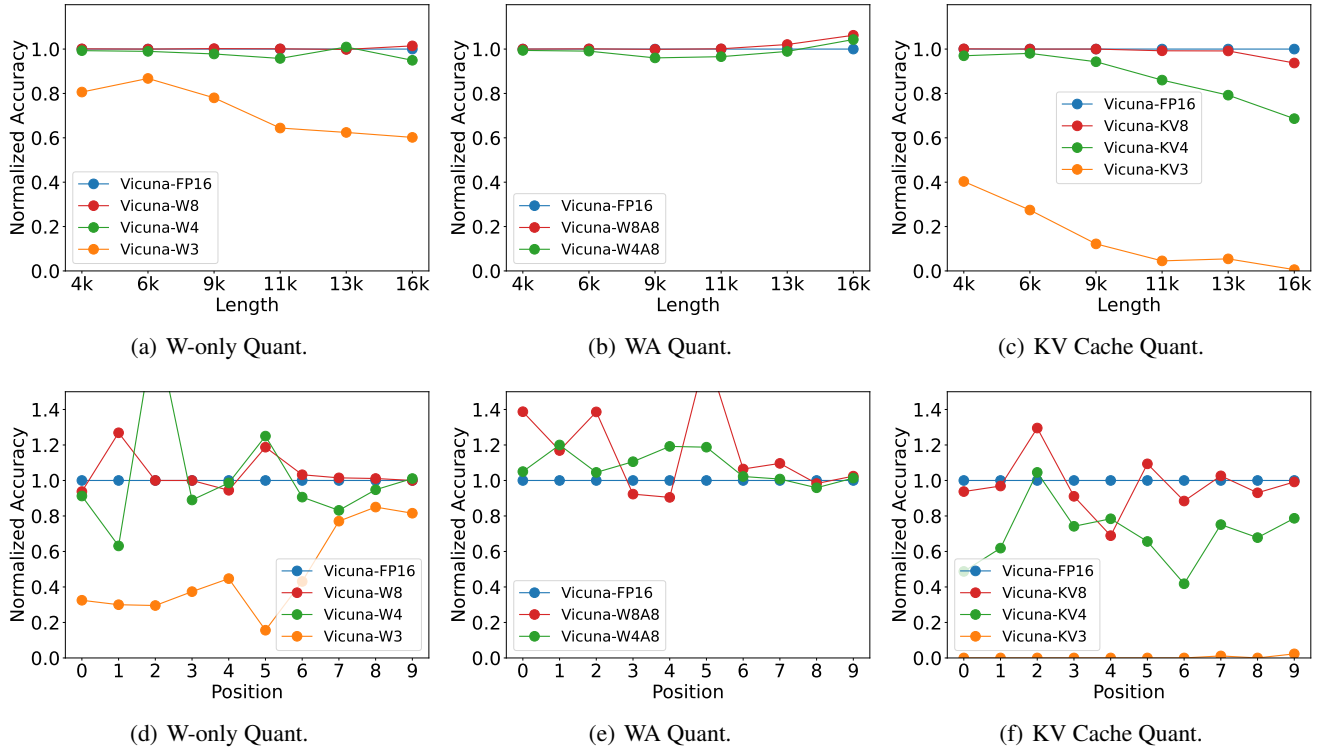


Figure 29: LongEval dataset. The effect of quantization on effective context length (a, b, c) and context position (d, e, f) for the Vicuna family.

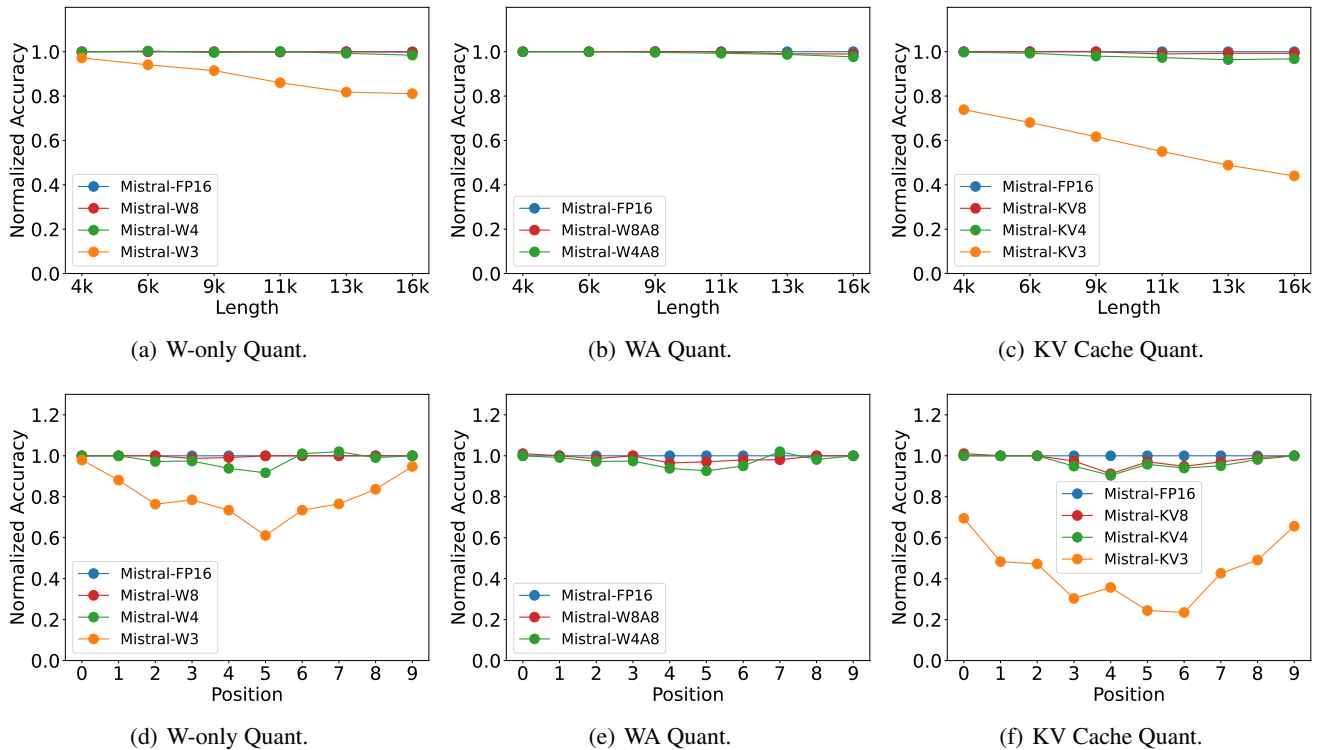


Figure 30: LongEval dataset. The effect of quantization on effective context length (a, b, c) and context position (d, e, f) for the Mistral family.

F.2. Introduction of Metrics

For the LongEval Dataset, the author extracts the last digital number from the output results of LLMs and compares it with the ground-truth value. If the two values match, it is considered correct; otherwise, it is deemed incorrect.

For the Multi-Document Question Answering dataset, we also use accuracy as the evaluation metric to judge whether any of the correct answers appear in the predicted output. Specifically, we normalize the results by removing articles and punctuations from the generated results and converting all letters to lowercase. Then, if the correct answers appear in the normalized results, it indicates a correct answer; otherwise, it is considered incorrect.




Figure 31: LongEval. The effect of quantization on effective context length (a, b, c) and context position (d, e, f) for ChatGLM2-6B.

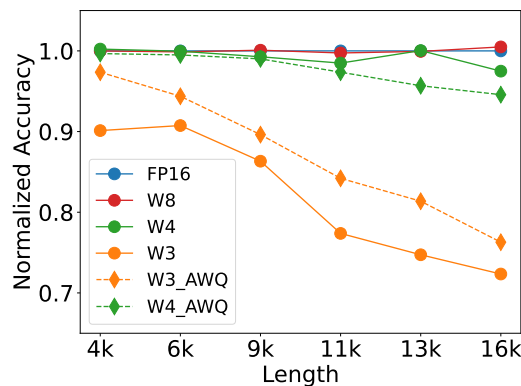


Figure 32: The effect of SOTA quantization methods on the LongEval dataset.

F.3. Additional Results on LongEval dataset

The LongChat family (16K) (Li et al., 2023), which is based on the LLaMA family, is very sensitive to KV Cache Quantization. We can not even quantize LongChat models to KV8 for the texts with more than 6K tokens, as shown in Figure 28 (c). The Weight-only Quantization is better than KV Cache Quantization, and we can quantize to W8 without




Figure 33: Multi-Doc QA. The effect of quantization on context position.

notable performance loss. The Weight-Activation Quantization is also very sensitive, and even W8A8 quantization has notable performance loss for the texts with more than 9K tokens. As for the effects of quantization on the text positions, for both Weight-only and KV Cache Quantization, lower bit-width has a slightly greater impact on the latter part of the text. Since the model itself exhibits relatively weaker performance on the earlier part of the text, we cannot assert that quantization has different effects on different positions of the texts based solely on the magnitude of performance degradation. As a consequence, the LLaMA-based LongChat family is more sensitive to quantization than other model families. Additionally, in long-context tasks, the models are more sensitive to KV Cache Quantization compared to Weight-only Quantization is also valid for the LongChat model family.

The Vicuna family (16K) (Zheng et al., 2023b), which is based on the LLaMA2 family, is also very sensitive to KV Cache Quantization. However, we can quantize Vicuna models to KV8 without significant accuracy loss. For the Weight-only and Weight-Activation Quantization, the W4 and W4A8 Vicuna models are good trade-offs for better efficiency and low-performance loss, as shown in Figure 29. As for the text positions, the conclusion is similar to that of the LongChat family. In addition, the phenomena of the **ChatGLM2-6B (32K) (Du et al., 2022)** are similar to that of Vicuna and can also be quantized to W4, W4A8, and KV4.

The Mistral family (32K) (Jiang et al., 2023) has better performance than other model families. The performance of the FP16 baselines from the Mistral family is nearly 100%. Although the Mistral family is still more sensitive to KV Cache Quantization than Weight-only Quantization, we can further quantize the KV Cache to 4-bit without significant loss. Another different phenomenon, which is also mentioned in Sec. 7, is that the middle of the texts is more sensitive to Weight-only Quantization and KV Cache Quantization, as shown in Figure 30.

The effects of SOTA quantization methods are shown in Figure 32. Specifically, we plot the performance curves under different bit-width by averaging the normalized performance on all the LLMs based on the performance of the FP16 LLMs. In this case, AWQ can even degrade the performance of W4 quantized LLMs, especially for the Vicuna family. While AWQ brings significant performance improvement for W3 quantized LLMs, it struggles to fully recover the performance to the level of FP16 due to the substantial accuracy loss incurred by W3 quantization, exceeding 20%.

F.4. Additional Results on Multi-document QA dataset

On all the model families we evaluated, we observe the U-shape performance curve as introduced by Liu et al. (2023a). For the Vicuna, ChatGLM, and Mistral families, we observe similar results as described in Appendix F.3. Generally speaking, in most cases, quantization has consistent effects across different positions, with only a few exceptions. For example, the beginning of the texts is more sensitive to Weight-only Quantization for Vicuna-13B.