

Exploring the Trade-Offs: Quantization Methods, Task Difficulty, and Model Size in Large Language Models From Edge to Giant

Jemin Lee¹, Sihyeong Park², Jinse Kwon¹, Jihun Oh³ and Yongin Kwon^{1*}

¹Electronics and Telecommunications Research Institute

²Korea Electronics Technology Institute

³Neubla

{leejaymin,kwonse,yongin.kwon}@etri.re.kr

sihyeong@keti.re.kr, oj9040@gmail.com

Abstract

Quantization has gained attention as a promising solution for the cost-effective deployment of large and small language models. However, most prior work has been limited to perplexity or basic knowledge tasks and lacks a comprehensive evaluation of recent models like Llama-3.3. In this paper, we conduct a comprehensive evaluation of instruction-tuned models spanning 1B to 405B parameters, applying four quantization methods across 13 datasets. Our findings reveal that (1) quantized models generally surpass smaller FP16 baselines, yet they often struggle with instruction-following and hallucination detection; (2) FP8 consistently emerges as the most robust option across tasks, and AWQ tends to outperform GPTQ in weight-only quantization; (3) smaller models can suffer severe accuracy drops at 4-bit quantization, while 70B-scale models maintain stable performance; (4) notably, *hard* tasks do not always experience the largest accuracy losses, indicating that quantization magnifies a model’s inherent weaknesses rather than simply correlating with task difficulty; and (5) an LLM-based judge (MT-Bench) highlights significant performance declines in Coding and STEM tasks, though it occasionally reports improvements in reasoning.

1 Introduction

Despite the remarkable performance of recent large and small language models (LLMs and SLMs), deploying them in resource-constrained environments remains challenging. Even models like Llama-3.3-70B (released in December 2024) and Llama-3.2-1B (released in September 2024) still involve billions of parameters, making them costly to run in both server and mobile-edge scenarios. Low-bit quantization has emerged as a popular solution to reduce the memory and computational overhead of these models. In particular, Post-Training Quantization (PTQ) [Frantar *et al.*, 2022; Lin *et al.*, 2024; Xiao *et al.*, 2023; Micikevicius *et al.*,

2022] is widely adopted, as Quantization Aware Training (QAT) often requires extensive retraining [Zhu *et al.*, 2023; Wan *et al.*, 2023].

However, existing research on quantization has largely relied on perplexity-based metrics [Frantar *et al.*, 2022; Lin *et al.*, 2024; Xiao *et al.*, 2023] and older benchmarks (e.g., ARC, HellaSwag, Winogrande, MMLU) [Yao *et al.*, 2023; Dettmers and Zettlemoyer, 2023; Liu *et al.*, 2023; Jin *et al.*, 2024; Li *et al.*, 2024; Dutta *et al.*, 2024], which have become too easy for current models and risk data contamination in recently trained LLMs. Moreover, more recent architectures like Llama-3.3, Llama-3.2, and Llama-3.1 have not been thoroughly investigated. This gap includes extreme scales, from 1B to over 405B parameters, and omits detailed category-level analysis using LLM-as-judge evaluation methods. Additionally, there has been limited manual inspection to refine and confirm evaluation results [Li *et al.*, 2024; Dutta *et al.*, 2024].

In this paper, we present a comprehensive evaluation of how quantization affects instruction-tuned LLMs. In particular, we aim to address the following research questions (RQs): **(RQ1)** Do quantized LLMs outperform smaller original models in most benchmarks, and how do they perform across diverse architectures and small language models (SLMs)? **(RQ2)** How do different quantization methods influence performance across a broad range of tasks, and are there significant differences in how specific approaches (e.g., GPTQ, AWQ, SmoothQuant, FP8) affect task accuracy? **(RQ3)** In what ways do model size and architecture affect the accuracy of quantized models? **(RQ4)** Does higher task difficulty necessarily correlate with greater accuracy degradation under quantization? and **(RQ5)** How does quantization impact the free-form conversation quality of LLMs when evaluated using the MT-Bench framework, which relies on LLMs as judges?

We perform our evaluation in a multi-cluster GPU environment, as illustrated in Figure 1. This setup consists of four servers, each with a distinct GPU configuration, and ensures consistent measurement conditions across all experiments (details in Appendix D).

Our study applies four quantization methods—GPTQ [Frantar *et al.*, 2022], AWQ [Lin *et al.*, 2024], SmoothQuant [Xiao *et al.*, 2023], and FP8 [Micikevicius *et al.*, 2022]—to instruction-tuned models ranging from 1B to 405B param-

*Corresponding author

Appendix available at: <https://arxiv.org/abs/2409.11055>

Model	Method	W/A	Storage (GB)	OpenLLM Leaderboard-v1 \uparrow							Avg.
				ARC-c (25-shot) acc_norm	GSM8k (5-shot) acc	HellaSwag (10-shot) acc_norm	MMLU (5-shot) acc	TruthfulQA (0-shot) mc2	Winogrande (5-shot) acc		
Llama-2-7B-Chat	FP16	16 / 16	14	53.16	21.91	78.92	47.24	78.92	45.32	72.13	53.11
	GPTQ*	4 / 16	3.5	51.28 ($\downarrow 1.88$)	13.87 ($\downarrow 8.04$)	72.17 ($\downarrow 6.75$)	43.10 ($\downarrow 4.14$)	44.12 ($\downarrow 1.20$)	71.27 ($\downarrow 0.86$)	49.30 ($\downarrow 3.81$)	
	AWQ	4 / 16	3.5	52.47 ($\downarrow 0.69$)	19.63 ($\downarrow 2.28$)	78.13 ($\downarrow 0.79$)	45.34 ($\downarrow 1.90$)	44.28 ($\downarrow 1.04$)	71.19 ($\downarrow 0.94$)	51.84 ($\downarrow 1.27$)	
Llama-2-13B-Chat	FP16	16 / 16	26	58.87	35.55	82.45	53.55	43.95	75.29	58.28	
	GPTQ*	4 / 16	6.5	57.51 ($\downarrow 1.36$)	32.22 ($\downarrow 3.33$)	81.35 ($\downarrow 1.10$)	52.36 ($\downarrow 1.19$)	41.74 ($\downarrow 2.21$)	75.76 ($\uparrow 0.47$)	56.82 ($\downarrow 1.46$)	
	AWQ	4 / 16	6.5	57.94 ($\downarrow 0.93$)	34.79 ($\downarrow 0.76$)	81.58 ($\downarrow 0.87$)	53.76 ($\uparrow 0.21$)	43.64 ($\downarrow 0.31$)	74.90 ($\downarrow 0.39$)	57.77 ($\downarrow 0.51$)	
Llama-2-70B-Chat	FP16	16 / 16	140	66.30	50.64	85.61	63.18	52.76	80.50	66.50	
	GPTQ*	4 / 16	35	62.88 ($\downarrow 3.42$)	50.27 ($\downarrow 0.37$)	84.98 ($\downarrow 0.63$)	61.57 ($\downarrow 1.61$)	51.13 ($\downarrow 1.63$)	79.32 ($\downarrow 1.18$)	65.03 ($\downarrow 1.47$)	
	AWQ	4 / 16	35	65.27 ($\downarrow 1.03$)	48.14 ($\downarrow 2.50$)	85.29 ($\downarrow 0.32$)	62.65 ($\downarrow 0.53$)	52.75 ($\downarrow 0.01$)	79.87 ($\downarrow 0.63$)	65.66 ($\downarrow 0.84$)	
Llama-3.1-8B-it	FP16	16 / 16	16	60.24	76.65	80.21	68.10	54.03	76.16	69.23	
	FP8	8 / 8	8	61.52 ($\uparrow 1.28$)	74.75 ($\downarrow 1.90$)	80.12 ($\downarrow 0.09$)	68.52 ($\uparrow 0.42$)	53.81 ($\downarrow 0.22$)	77.43 ($\uparrow 1.27$)	69.36 ($\uparrow 0.13$)	
	GPTQ*	4 / 16	4	61.43 ($\uparrow 1.19$)	72.33 ($\downarrow 4.32$)	78.36 ($\downarrow 1.85$)	66.85 ($\downarrow 1.25$)	53.60 ($\downarrow 0.43$)	75.22 ($\downarrow 0.94$)	67.97 ($\downarrow 1.26$)	
	GPTQ**	4 / 16	4	59.81 ($\downarrow 0.43$)	69.98 ($\downarrow 6.67$)	78.53 ($\downarrow 1.68$)	66.07 ($\downarrow 2.03$)	50.45 ($\downarrow 3.58$)	76.64 ($\uparrow 0.48$)	66.91 ($\downarrow 2.32$)	
	GPTQ**	8 / 16	8	61.01 ($\uparrow 0.77$)	75.81 ($\downarrow 0.84$)	80.27 ($\downarrow 0.06$)	68.21 ($\uparrow 0.11$)	54.03 (0.00)	77.19 ($\uparrow 1.03$)	69.42 ($\uparrow 0.19$)	
	SmoothQuant	8 / 8	8	60.75 ($\uparrow 0.51$)	76.12 ($\downarrow 0.53$)	80.08 ($\downarrow 0.13$)	68.22 ($\uparrow 0.12$)	53.85 ($\downarrow 0.18$)	77.11 ($\uparrow 0.95$)	69.36 ($\uparrow 0.13$)	
	AWQ	4 / 16	4	58.53 ($\downarrow 1.71$)	73.39 ($\downarrow 3.26$)	79.10 ($\downarrow 1.11$)	66.26 ($\downarrow 1.84$)	51.87 ($\downarrow 2.16$)	75.37 ($\downarrow 0.79$)	67.42 ($\downarrow 1.81$)	
Llama-3.1-70B-it	FP16	16 / 16	140	69.54	88.70	86.74	82.30	59.85	85.40	78.76	
	FP8	8 / 8	70	69.45 ($\downarrow 0.09$)	88.25 ($\downarrow 0.45$)	86.69 ($\downarrow 0.05$)	82.02 ($\downarrow 0.28$)	59.80 ($\downarrow 0.05$)	85.08 ($\downarrow 0.32$)	78.55 ($\downarrow 0.21$)	
	GPTQ*	4 / 16	35	69.80 ($\uparrow 0.26$)	89.54 ($\uparrow 0.84$)	86.28 ($\downarrow 0.46$)	81.40 ($\downarrow 0.90$)	59.37 ($\downarrow 0.48$)	84.69 ($\downarrow 0.71$)	78.51 ($\downarrow 0.25$)	
	GPTQ**	4 / 16	35	69.97 ($\uparrow 0.43$)	89.76 ($\uparrow 1.06$)	86.26 ($\downarrow 0.48$)	81.97 ($\downarrow 0.33$)	58.74 ($\downarrow 1.11$)	84.53 ($\downarrow 0.87$)	78.54 ($\downarrow 0.22$)	
	GPTQ**	8 / 16	70	69.03 ($\downarrow 0.51$)	87.95 ($\downarrow 0.75$)	86.29 ($\downarrow 0.45$)	82.17 ($\downarrow 0.13$)	58.94 ($\downarrow 0.91$)	84.53 ($\downarrow 0.87$)	78.15 ($\downarrow 0.61$)	
	SmoothQuant	8 / 8	70	70.05 ($\uparrow 0.51$)	88.55 ($\downarrow 0.15$)	86.56 ($\downarrow 0.18$)	82.10 ($\downarrow 0.20$)	60.39 ($\uparrow 0.54$)	85.24 ($\downarrow 0.16$)	78.82 ($\uparrow 0.06$)	
	AWQ	4 / 16	35	69.80 ($\uparrow 0.26$)	90.83 ($\uparrow 2.13$)	86.18 ($\downarrow 0.56$)	81.33 ($\downarrow 0.97$)	59.68 ($\downarrow 0.17$)	84.37 ($\downarrow 1.03$)	78.70 ($\downarrow 0.06$)	
Llama-3.1-405B-it	FP16	16 / 16	810	73.72	94.84	88.40	83.98	65.42	85.00	81.89	
	FP8	8 / 8	405	73.12 ($\downarrow 0.60$)	95.38 ($\uparrow 0.54$)	88.32 ($\downarrow 0.08$)	85.91 ($\uparrow 1.93$)	64.79 ($\downarrow 0.63$)	85.63 ($\uparrow 0.63$)	82.19 ($\uparrow 0.30$)	
	GPTQ**	4 / 16	202.5	72.10 ($\downarrow 1.62$)	94.24 ($\downarrow 0.60$)	88.17 ($\downarrow 0.23$)	85.79 ($\uparrow 1.81$)	64.80 ($\downarrow 0.62$)	85.48 ($\uparrow 0.48$)	81.76 ($\downarrow 0.13$)	
	SmoothQuant	8 / 8	405	72.01 ($\downarrow 1.71$)	92.72 ($\downarrow 2.12$)	87.53 ($\downarrow 0.87$)	73.28 ($\downarrow 10.70$)	65.19 ($\downarrow 0.23$)	85.95 ($\uparrow 0.95$)	79.45 ($\downarrow 2.44$)	
	AWQ	4 / 16	202.5	73.98 ($\uparrow 0.26$)	94.84 (0.00)	88.04 ($\downarrow 0.36$)	85.71 ($\uparrow 1.73$)	64.25 ($\downarrow 1.17$)	86.35 ($\uparrow 1.35$)	82.20 ($\uparrow 0.31$)	
Llama-3.2-1B-it	FP16	16 / 16	2	42.06	33.36	59.62	45.44	43.82	62.59	47.81	
	FP8	8 / 8	1	41.98 ($\downarrow 0.08$)	34.04 ($\uparrow 0.68$)	59.20 ($\downarrow 0.42$)	45.38 ($\downarrow 0.06$)	43.18 ($\downarrow 0.64$)	62.67 ($\uparrow 0.08$)	47.74 ($\downarrow 0.07$)	
	GPTQ***	4 / 16	0.5	34.90 ($\downarrow 7.16$)	8.04 ($\downarrow 25.32$)	43.14 ($\downarrow 16.48$)	40.20 ($\downarrow 5.24$)	41.92 ($\downarrow 1.90$)	57.85 ($\downarrow 4.74$)	37.67 ($\downarrow 10.14$)	
	SmoothQuant	8 / 8	0.5	42.06 (0.00)	33.13 ($\downarrow 0.23$)	59.51 ($\downarrow 0.11$)	45.19 ($\downarrow 0.25$)	43.40 ($\downarrow 0.42$)	61.48 ($\downarrow 1.11$)	47.46 ($\downarrow 0.35$)	
	AWQ	4 / 16	0.5	38.48 ($\downarrow 3.58$)	17.97 ($\downarrow 15.39$)	53.82 ($\downarrow 5.80$)	40.51 ($\downarrow 4.39$)	42.61 ($\downarrow 1.21$)	59.12 ($\downarrow 3.47$)	42.09 ($\downarrow 5.72$)	
Llama-3.2-3B-it	FP16	16 / 16	6	52.13	64.52	73.08	59.59	47.79	69.61	61.45	
	FP8	8 / 8	3	51.37 ($\downarrow 0.76$)	63.31 ($\downarrow 1.21$)	73.04 ($\downarrow 0.04$)	59.74 ($\uparrow 0.15$)	49.98 ($\uparrow 0.19$)	69.14 ($\downarrow 0.47$)	61.10 ($\downarrow 0.35$)	
	GPTQ***	4 / 16	1.5	50.26 ($\downarrow 1.87$)	60.20 ($\downarrow 4.32$)	71.19 ($\downarrow 1.89$)	57.90 ($\downarrow 1.69$)	49.52 ($\downarrow 0.27$)	68.75 ($\downarrow 0.86$)	59.64 ($\downarrow 1.81$)	
	SmoothQuant	8 / 8	1.5	51.37 ($\downarrow 0.76$)	63.76 ($\downarrow 0.76$)	72.61 ($\downarrow 0.47$)	56.69 ($\uparrow 0.10$)	49.72 ($\downarrow 0.07$)	69.61 (0.00)	61.13 ($\downarrow 0.32$)	
Llama-3.3-70B-it	FP16	16 / 16	140	71.67	90.83	86.39	82.20	60.90	83.98	79.33	
	GPTQ***	4 / 16	35	69.71 ($\downarrow 1.96$)	89.39 ($\downarrow 1.44$)	85.58 ($\downarrow 0.81$)	81.63 ($\downarrow 0.57$)	61.25 ($\uparrow 0.35$)	84.21 ($\uparrow 0.23$)	78.63 ($\downarrow 0.70$)	
	AWQ	4 / 16	35	70.82 ($\downarrow 0.85$)	88.17 ($\downarrow 2.66$)	85.73 ($\downarrow 0.66$)	81.45 ($\downarrow 0.75$)	60.82 ($\downarrow 0.08$)	83.98 (0.00)	78.50 ($\downarrow 0.83$)	

Table 1: Evaluation of Llama families on OpenLLM Leaderboard-v1. *, **, and *** denote the use of AutoGPTQ, llmcompressor, and AutoRound for GPTQ quantization, respectively.

of model quantization on the performance of LLMs, focusing on various aspects. For instance, Yao et al. [Yao *et al.*, 2023] investigated the impact of quantization on both weights and activations in language modeling tasks. In contrast, Liu et al. [Liu *et al.*, 2023] concentrated solely on evaluating three emergent abilities of quantized LLMs, neglecting crucial tasks such as trustworthiness, dialogue, and long-context processing.

Hong et al. [Hong *et al.*, 2024] expanded the scope by examining trustworthiness dimensions in the assessment of LLM compression techniques. However, most studies have predominantly relied on accuracy as the primary evaluation metric, with limited attention paid to alternative metrics. For example, Zhang et al. [Zhang *et al.*, 2024] proposed additional evaluation metrics, including fluency, informativeness, coherence, and harmlessness, alongside accuracy.

Efforts to establish more comprehensive evaluation benchmarks have also been made. Jaiswal et al. [Jaiswal *et al.*, 2023]

developed a benchmark from existing datasets to evaluate compressed models, while Li et al. [Li *et al.*, 2024] and Jin et al. [Jin *et al.*, 2024] assessed various quantization techniques across different tasks. Namburi et al. [Namburi *et al.*, 2023] explored how compression and pruning affect the parametric knowledge of LLMs.

Dutta et al. [Dutta *et al.*, 2024] proposed a new metric from the perspective of “Flip” errors, emphasizing the importance of evaluating model robustness by accounting for inconsistencies and reversals in predictions, thereby going beyond traditional accuracy-focused metrics. Kurtic et al. [Kurtic *et al.*, 2024] investigated the accuracy-performance trade-offs in quantizing LLMs, evaluating formats like FP8, INT8, and INT4 across various tasks and proposing practical guidelines for efficient LLM deployment at different model scales.

Xu et al. [Xu *et al.*, 2024] explored the challenges of evaluating multilingual LLMs across diverse languages and cultures,

Model	Method	W/A	Storage (GB)	OpenLLM Leaderboard-v2 \uparrow							Avg.
				BBH (3-shot)	GPQA (0-shot)	IFEval (0-shot)	Math-Lvl-5 (4-shot)	MMLU-PRO (5-shot)	MuSR (0-shot)		
Llama-2-7B-Chat	FP16	16 / 16	14	12.23	1.59	35.31	1.93	11.04	8.89	11.83	
	GPTQ*	4 / 16	3.5	7.96 (↓4.27)	0.85 (↓0.74)	30.59 (↓4.72)	1.43 (↓0.50)	7.58 (↓3.46)	3.73 (↓5.16)	8.69 (↓3.14)	
	AWQ	4 / 16	3.5	10.78 (↓1.45)	3.54 (↑1.95)	31.57 (↓3.74)	1.81 (↓0.12)	7.75 (↓3.29)	10.62 (↑1.73)	11.01 (↓0.82)	
Llama-2-13B-Chat	FP16	16 / 16	26	16.87	4.03	37.45	1.56	15.36	10.00	14.21	
	GPTQ*	4 / 16	6.5	16.92 (↑0.05)	4.27 (↑0.24)	33.40 (↓4.05)	2.21 (↑0.65)	15.44 (↑0.08)	8.30 (↓1.70)	13.42 (↓0.79)	
	AWQ	4 / 16	6.5	16.54 (↓0.33)	6.96 (↑2.93)	33.28 (↓4.17)	1.80 (↑0.24)	15.57 (↑0.21)	10.39 (↑0.39)	14.09 (↓0.12)	
Llama-2-70B-Chat	FP16	16 / 16	140	29.42	6.72	44.11	2.64	25.01	6.32	19.04	
	GPTQ*	4 / 16	35	25.80 (↓3.62)	5.25 (↓1.47)	41.52 (↓2.59)	3.27 (↑0.63)	23.87 (↓1.14)	7.06 (↑0.74)	17.80 (↓1.24)	
	AWQ	4 / 16	35	28.63 (↓0.79)	6.72 (0.00)	42.89 (↓1.22)	1.95 (↓0.69)	24.44 (↓0.57)	5.56 (↓0.76)	18.37 (↓0.67)	
Llama-3.1-8B-it	FP16	16 / 16	16	30.11	6.23	50.09	11.69	30.90	8.88	22.98	
	FP8	8 / 8	8	29.20 (↓0.91)	5.49 (↓0.74)	49.16 (↓0.93)	12.01 (↑0.32)	30.92 (↑0.02)	6.95 (↓1.93)	22.29 (↓0.69)	
	GPTQ*	4 / 16	4	25.86 (↓4.25)	7.20 (↑0.97)	47.95 (↓2.14)	9.49 (↓2.20)	29.60 (↓1.30)	6.03 (↓2.85)	21.02 (↓1.96)	
	GPTQ**	4 / 16	4	25.83 (↓4.28)	6.72 (↑0.49)	44.81 (↓5.28)	8.85 (↓2.84)	28.16 (↓2.74)	10.58 (↑1.70)	20.83 (↓2.15)	
	GPTQ**	8 / 16	8	29.97 (↓0.14)	6.23 (0.00)	50.53 (↑0.44)	11.94 (↑0.25)	31.19 (↑0.29)	7.80 (↓1.08)	22.94 (↓0.04)	
	SmoothQuant	8 / 8	8	30.19 (↑0.08)	2.56 (↓3.67)	50.25 (↑0.16)	12.77 (↑1.08)	30.75 (↓0.15)	8.12 (↓0.76)	22.44 (↓0.54)	
Llama-3.1-70B-it	FP16	16 / 16	140	55.90	16.48	75.48	28.68	48.00	19.32	40.64	
	FP8	8 / 8	70	55.54 (↓0.36)	16.24 (↓0.24)	75.75 (↑0.27)	28.92 (↑0.24)	47.84 (↓0.16)	19.35 (↑0.03)	40.61 (↓0.03)	
	GPTQ*	4 / 16	35	53.65 (↓2.25)	17.70 (↑1.22)	73.26 (↓2.22)	27.26 (↓1.42)	47.49 (↓0.51)	20.33 (↑1.01)	39.95 (↓0.69)	
	GPTQ**	4 / 16	35	55.79 (↓0.11)	14.04 (↓2.44)	72.71 (↓2.77)	26.16 (↓2.52)	46.97 (↓1.03)	16.93 (↓2.39)	38.77 (↓1.87)	
	GPTQ**	8 / 16	70	54.79 (↓1.11)	2.81 (↓13.67)	66.66 (↓8.82)	29.06 (↑0.38)	47.56 (↓0.44)	20.42 (↑1.10)	36.88 (↓3.76)	
	SmoothQuant	8 / 8	70	55.06 (↓0.84)	16.24 (↓0.24)	74.78 (↓0.70)	27.90 (↓0.78)	47.20 (↓0.80)	20.32 (↑1.00)	40.25 (↓0.39)	
Llama-3.1-405B-it	FP16	16 / 16	810	66.81	26.25	76.18	37.06	60.01	19.86	47.70	
	FP8	8 / 8	405	65.22 (↓1.59)	27.37 (↑1.22)	72.44 (↓3.74)	35.86 (↓1.20)	59.67 (↓0.34)	17.83 (↓2.03)	46.42 (↓1.28)	
	GPTQ**	4 / 16	202.5	66.21 (↓0.60)	23.57 (↓2.68)	72.90 (↓3.28)	34.74 (↓2.32)	59.11 (↓0.90)	18.92 (↓0.94)	45.91 (↓1.79)	
	SmoothQuant	8 / 8	405	54.46 (↓12.35)	16.73 (↓9.52)	70.34 (↓5.84)	35.01 (↓2.05)	18.24 (↓11.77)	18.60 (↓1.26)	35.56 (↓12.14)	
	AWQ	4 / 16	202.5	65.50 (↓1.31)	26.50 (↑0.25)	47.52 (↓28.66)	38.13 (↑1.07)	58.63 (↓1.38)	19.69 (↓0.17)	42.66 (↓5.04)	
	Llama-3.2-1B-it	FP16	16 / 16	2	8.32	1.34	41.61	4.20	10.60	3.70	11.63
FP8		8 / 8	1	8.83 (↑0.51)	1.59 (↑0.25)	43.18 (↑1.57)	3.51 (↓0.69)	10.28 (↓0.32)	3.49 (↓0.21)	11.81 (↑0.18)	
GPTQ***		4 / 16	0.5	3.98 (↓4.34)	0.85 (↓0.49)	25.60 (↓16.01)	0.30 (↓3.90)	6.72 (↓3.88)	1.86 (↓1.84)	6.55 (↓5.08)	
SmoothQuant		8 / 8	0.5	8.22 (↓0.10)	4.03 (↑2.69)	41.68 (↑0.07)	3.25 (↓0.95)	10.17 (↓0.43)	3.27 (↓0.43)	11.77 (↑0.14)	
Llama-3.2-3B-it	FP16	16 / 16	6	20.77	7.69	53.56	11.20	22.24	6.92	20.40	
	FP8	8 / 8	3	21.00 (↑0.23)	7.69 (0.00)	53.32 (↓0.24)	9.78 (↓1.42)	21.99 (↓0.25)	7.30 (↑0.38)	20.18 (↓0.22)	
	GPTQ***	4 / 16	1.5	18.82 (↓1.95)	5.98 (↓1.71)	52.81 (↓0.75)	8.05 (↓3.14)	19.31 (↓2.93)	4.19 (↓2.73)	18.19 (↓2.21)	
	SmoothQuant	8 / 8	1.5	21.01 (↑0.24)	9.65 (↓1.96)	53.96 (↑0.40)	9.58 (↓1.62)	21.83 (↓0.41)	6.07 (↓0.85)	20.35 (↓0.05)	
Llama-3.3-70B-it	FP16	16 / 16	140	56.78	30.40	69.03	30.95	49.78	22.28	43.20	
	GPTQ***	4 / 16	35	52.48 (↓4.31)	28.94 (↓1.46)	65.71 (↓3.32)	28.55 (↓2.40)	48.24 (↓1.54)	21.38 (↓0.90)	40.88 (↓2.32)	
	AWQ	4 / 16	35	55.89 (↓0.89)	28.69 (↓1.71)	69.61 (↑0.58)	29.42 (↓1.53)	48.57 (↓1.21)	20.67 (↓1.61)	42.14 (↓1.06)	

Table 2: Evaluation of Llama families on OpenLLM Leaderboard-v2. *, **, and *** denote the use of AutoGPTQ, llmcompressor, and AutoRound for GPTQ quantization, respectively.

emphasizing the development of culturally and linguistically inclusive benchmarks for fair evaluation. Liu et al. [Liu et al., 2024] provided a comprehensive examination of the generalization ability, focusing on their performance across various tasks and datasets. This study highlights the importance of designing benchmarks and metrics that accurately reflect real-world applications while identifying the limitations of current evaluation strategies.

To our knowledge, no prior study has comprehensively examined the effects of quantization across a wide range of model sizes—from 1B to 405B parameters—encompassing both SLMs and LLMs, including the latest architectures such as Llama-3.1, Llama-3.2, and Llama-3.3. Furthermore, existing research has not conducted detailed category-level analyses through cross-architecture comparisons or employed manual inspection using LLM-as-judge qualitative methods. Additionally, no work has compared the trends observed in LLM-as-judge (MT-Bench) evaluations with leaderboard results to

identify new trends in quantization impacts.

3 Evaluation Procedure

To handle LLMs, which cannot be processed on a single server, and to ensure fast and reliable evaluations, we developed a structured evaluation pipeline based on a multi-node cluster setup. Figure 1 presents an overview of the implemented pipeline for evaluating quantized LLMs. The evaluated LLMs include the Vicuna, Gemma, and Llama families, ranging in size from 1B to 405B. Each model is quantized using GPTQ, AWQ, SmoothQuant, and FP8 methods. The evaluation is conducted using lm-eval and MT-Bench as benchmarking tools. The multi-node cluster used for evaluation is implemented with vLLM and consists of four servers: H100-80Gx8, A100-80Gx4, RTX 6000-48Gx4, and A6000-48Gx4. Additionally, the Huggingface library is integrated into the pipeline to support model hosting and benchmarking. The evaluation is distributed across a multi-cluster environment to ensure a

thorough performance assessment. If vLLM cannot be used for processing, we used the Huggingface Accelerate library instead, which is slower but shows better comparability. The versions of all tools used are provided in Appendix D.1.

4 Experimental Setup

4.1 Datasets

We conducted a comprehensive evaluation of the quantized LLMs on widely adopted benchmarks, grouped into six main categories: CommonSenseQA (ARC, HellaSwag, Winogrande), Knowledge and Language Understanding (MMLU, GPQA, MMLU-PRO, BBH, MuSR), Instruction Following (IFEval), Hallucination (TruthfulQA), Mathematics (GSM8K, MATH-Lvl-5), and Dialogue (MT-Bench). Additional information about these datasets can be found in Appendix A, and an overview of all benchmarks is provided in Table 3 (Appendix).

4.2 Reproducibility Details

Both OpenLLM Leaderboard V1 and V2 follow the same methodology outlined on the HuggingFace Leaderboard’s page, including identical normalization procedures. Additional information on leaderboard calculations is provided in Appendix D.2.

We used the *greedy decoding* strategy to maintain deterministic output tokens across runs. The detailed configuration for the *greedy decoding* strategy is presented in Appendix D.1, which also lists the specific versions of each package used. Furthermore, we record all random seeds for Python, NumPy, Torch, and few-shot setups in Appendix D.1, ensuring the reproducibility of our experimental results.

4.3 Quantization Methods and Calibration Data

We evaluated multiple PTQ methods, including GPTQ, AWQ, SmoothQuant, and FP8. GPTQ and AWQ focus on weight-only quantization, while SmoothQuant and FP8 apply to both weights and activations. For a detailed overview of each quantization method, refer to Appendix D.3, with configuration details and group sizes described in Appendix D.4.

The selection and configuration of calibration datasets are crucial for maintaining consistent performance across models. We used default settings for the number of samples and sequence lengths, as detailed in Appendix D.5 and summarized in Table 9. These settings may vary depending on the specific algorithms and tools but generally ensure stable results for our experiments.

4.4 Models

We applied quantization techniques to 12 instruction-tuned open LLMs, including the Vicuna [Zheng *et al.*, 2023], Gemma [Team *et al.*, 2024], and Llama [Dubey *et al.*, 2024] families, with model sizes ranging from 1B to 405B. These models were released between June 2023 and December 2024 and were downloaded from HuggingFace’s model sources.

All models were evaluated using 13 benchmark datasets, applying GPTQ, AWQ, SmoothQuant, and FP8 quantization. However, due to runtime limitations (over 30 days), we did not

measure the original model accuracy or test all GPTQ configurations for Llama-3.1-405B. Also, due to space limitations, the experimental results for the Vicuna and Gemma models are provided in the Appendix B

5 Experimental Results

This section presents experimental results addressing five research questions, detailing the performance impact of quantization across 13 datasets for three model families of varying sizes. Table 1 and Table 2 summarize the experimental results from OpenLLM Leaderboard-v1 and v2, respectively.

5.1 RQ1: Do quantized LLMs outperform smaller original models on most benchmarks, and how do they fare across different architectures and SLMs?

We observe that quantized LLMs generally outperform smaller, uncompressed models across a wide range of benchmarks. For instance, a 4-bit Llama-2-13B (6.5 GB) outperforms an FP16 Llama-2-7B (14 GB) on most tasks, despite its reduced size. However, in TruthfulQA (hallucination testing) and IFEval (instruction-following), the FP16 Llama-2-7B still performs better, indicating that quantization can compromise alignment and adherence to instructions.

Similarly, quantizing Llama-3.1-405B to 4 bits (202.5 GB) yields higher accuracy than the FP16 Llama-3.1-70B (140 GB) across various tasks; yet, the instruction-following IFEval benchmark again highlights a shortfall in the quantized model. This performance gap holds across different model architectures: although Llama-3.3-70B demonstrates improvements over Llama-3.1-70B, a 4-bit Llama-3.1-405B can still outperform the uncompressed Llama-3.3-70B.

In edge-focused SLMs, quantization produces significantly larger improvements. For example, quantizing Llama-3.2-3B (SmoothQuant) improves accuracy by 13.32% on OpenLLM Leaderboard-v1 and 8.72% on OpenLLM Leaderboard-v2 compared to the FP16-Llama-3.2-1B. Such improvements exceed the margin typically observed in larger models (7B to 405B).

RQ1 Findings: Quantized LLMs consistently outperform smaller models across most benchmarks and maintain this advantage across different architectures, with significant gains observed in both large models and edge-focused SLMs. However, tasks like instruction-following (IFEval) and hallucination detection (TruthfulQA) remain challenging for quantized models.

5.2 RQ2: How do different quantization methods affect the performance of models across diverse tasks? Are there noticeable differences in how specific methods (e.g., GPTQ, AWQ, SmoothQuant, FP8) impact task accuracy?

In most cases, weight-only quantization methods (GPTQ and AWQ) and activation quantization methods (SmoothQuant and FP8) exhibit similar performance. However, for Llama-3.1-405B, SmoothQuant’s activation quantization resulted in

a significant accuracy drop compared to other methods, with an average decrease of up to 10.86% compared to FP8 on the OpenLLM Leaderboard-v2 datasets. This decline occurs because SmoothQuant was originally designed to handle the high activation ranges observed in models up to the size of OPT-175B. Consequently, at the 405B scale of Llama-3.1, the algorithm likely did not account for certain factors that are critical at this larger scale. In contrast, for smaller models, such as SLMs with 1B and 3B parameters, the average accuracy drop remains below 1%.

When comparing weight-only quantization methods, AWQ consistently outperforms GPTQ across various LLMs on overall benchmark scores. Additionally, different implementations of GPTQ, such as AutoGPTQ and llmcompressor, demonstrate notable performance differences, with the oldest GPTQ implementation library, AutoGPTQ, still maintaining stable and consistent performance.

When both weight and activation quantization are required at 8 bits, FP8 proves to be highly effective, even on challenging tasks from the OpenLLM Leaderboard-v2. This effectiveness spans models of all sizes, from the largest, such as Llama-3.1-405B, to the smallest, like Llama-3.2-1B. Therefore, FP8 offers greater stability compared to SmoothQuant and is advantageous to use when supported by hardware such as the NVIDIA H100 GPU and RTX 6000 Ada, as it provides benefits in both latency and throughput. Additionally, applying FP8 to both weights and activations allows for a reduction of the KV cache size by half, which is highly beneficial during LLM decoding phases where I/O bottlenecks are a concern. This FP8 KV cache feature is supported by vLLM.

RQ2 Findings: FP8 is the most stable option across all model sizes and tasks, particularly in large LLMs where SmoothQuant performs poorly, whereas AWQ regularly outperforms GPTQ in weight-only quantization, and specialized hardware enhances FP8’s advantages.

5.3 RQ3: How do model size and architecture influence the accuracy of quantized models?

We evaluated GPTQ across 13 datasets and observed that its accuracy can degrade significantly under 4-bit quantization—a behavior not reflected in perplexity-based evaluations alone. For smaller models, such as Llama-3.2-1B, 4-bit quantization causes particularly severe accuracy drops (e.g., -25.32% on GSM8k and -16.01% on IFEval). These declines tend to be more pronounced in GPTQ than in AWQ, suggesting that SmoothQuant or FP8 at 8-bit may be necessary to maintain accuracy for 1B-scale models.

With mid-sized models like Llama-3.1-8B (GPTQ-8bit), we noted average accuracy improvements of +2.51% and +2.11% over 4-bit on OpenLLM Leaderboard-v1 and OpenLLM Leaderboard-v2, respectively. However, at the 70B scale, 4-bit outperformed 8-bit by +1.89%, indicating a reverse trend for larger models. To examine architectural differences at a fixed scale, we compared Llama-2, Llama-3.1, and Llama-3.3 at the 70B size and found that AWQ consistently surpassed GPTQ, delivering stable accuracy even at 4-bit quantization.

RQ3 Findings: In smaller LLMs, 4-bit quantization often leads to significant accuracy loss (especially with GPTQ), whereas 70B-scale models can maintain stable performance with 4-bit. Although size primarily drives quantization difficulty, architectural variations within the same scale can also influence accuracy. Nonetheless, AWQ consistently outperforms GPTQ across diverse tasks and model families.

5.4 RQ4: Does higher task difficulty always correlate with greater accuracy degradation under quantization?

Contrary to common assumptions, tasks widely considered challenging (e.g., GSM8K, MMLU, Math-Lvl-5) did not consistently show the greatest performance drops when quantized. For instance, among 12 tested models quantized with AWQ, MMLU accuracy remained largely unchanged, with some models even exhibiting slight improvements. In contrast, tasks generally regarded as less knowledge-intensive (ARC-challenge, TruthfulQA, Winogrande) occasionally experienced declines exceeding -3%.

Rather than strictly depending on task difficulty, quantization appears to magnify existing weaknesses in a model’s ability to handle specific forms of reasoning, such as commonsense or mathematical inference. As highlighted in RQ3, smaller models (2B–7B) are especially vulnerable to computational reasoning tasks like GSM8K, exhibiting sharper performance drops. Moreover, as noted in RQ2, different quantization methods can produce varying degrees of accuracy loss, making it difficult to draw definitive conclusions from a single task perspective.

RQ4 Findings: Difficult tasks do not always suffer the largest accuracy loss under quantization. The impact varies by model architecture and the chosen quantization method, leading some *hard* tasks to remain stable while *easier* tasks occasionally show bigger drops. In essence, quantization amplifies a model’s existing weaknesses, particularly in commonsense, logical, or mathematical reasoning.

5.5 RQ5: How does quantization impact the free-form conversation quality of LLMs when evaluated using the MT-Bench framework, which employs LLMs as judges?

Category-Level Analysis. Figure 2 presents a detailed breakdown of three models (Llama-3.1, Llama-3.2, and Llama-3.3) across MT-Bench categories. Quantized LLMs suffer the largest score degradation in *Coding* and *STEM*. For *Coding*, manual inspections reveal that GPT4 often assigns lower scores when the generated code contains fewer examples or insufficient comments. In *STEM* tasks, correctness matters when a definite solution exists, and concise logical explanations are important when an exact answer does not exist. Quantized models frequently either provided overly verbose justifications or produced incorrect statements, leading to lower scores. In

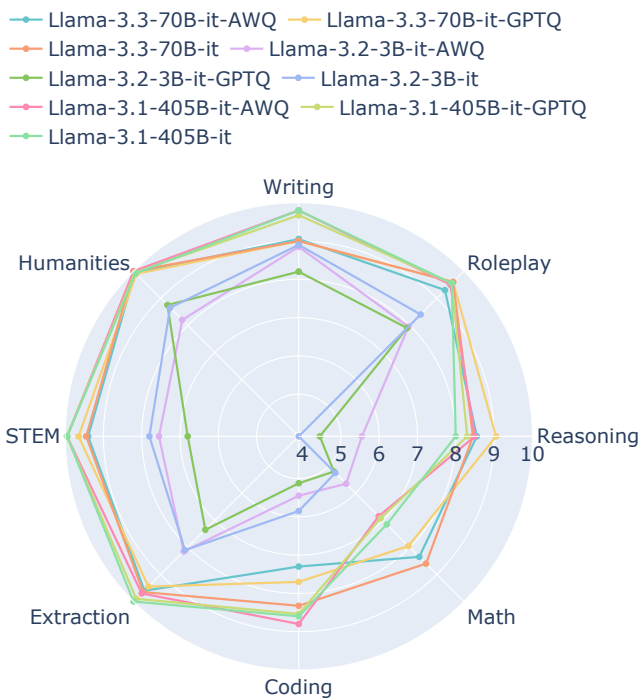


Figure 2: Category-wise MT-Bench scores of three quantized LLMs (Llama-3.3, Llama-3.2, and Llama-3.1) evaluated using AWQ and GPTQ methods. It highlights the performance differences across categories, including Writing, Roleplay, Reasoning, Math, Coding, Extraction, STEM, and Humanities, demonstrating the impact of quantization on diverse tasks.

contrast, the *Reasoning* category showed an increase of about 1 point with quantization. Manual checks reveal that concise answers tend to receive higher GPT4 scores, and quantized models often responded more concisely than their original counterparts. The detailed results of this manual inspection, along with the full text responses, can be found in Appendix C.4. Also, Table 6 in the Appendix C.1 lists category-wise MT-Bench scores for all models.

Limitations of GPT4-Based Evaluation. Consistent with the findings in MT-Bench, GPT4 sometimes misjudges incorrect responses as correct, particularly for math and reasoning tasks. Although reference-guided judging and chain-of-thought prompting can mitigate such errors [Zheng *et al.*, 2023], they do not eliminate them entirely. In *Reasoning* task, GPT4 erroneously considered a wrong answer correct, boosting quantized models’ scores. Conversely, there were *STEM* questions where both original and quantized models provided accurate answers, yet GPT4 mistakenly marked them as incorrect. **These misjudged cases are described in the Appendix C.5.**

For models like Llama-3.3-70B, categories such as *Humanities* consistently scored near or at the maximum (10 points). In these cases, quantized versions also achieved scores in the high 9-point range, making it difficult to discern meaningful quality differences through manual inspection. This is because it is difficult to clearly understand the reasoning behind the results by only looking at GPT4’s judgment statements, es-

pecially when the score differences are marginal, such as 1-2 points. Hence, for the latest large-scale models, more sensitive metrics or superior judging models may be required to evaluate subtle quality gaps.

Multi-Turn Analysis. Table 7 in the Appendix C.2 presents the average scores of multi-turn conversations across 12 models. Among smaller models (e.g., 2B, 7B, 8B), some exhibit slight score improvements; however, this trend is not consistent across all small-scale models. In contrast, larger models (e.g., 13B, 70B) generally experience score declines, although there are exceptions where AWQ enhances performance. Additionally, accuracy losses become more noticeable in the second turn of multi-turn interactions.

These observations indicate that establishing a clear trend based solely on model size or type is challenging, as the impact of quantization depends on a combination of factors, including model architecture, quantization method, and task complexity. Furthermore, when comparing quantization methods, AWQ typically outperforms GPTQ, which is consistent with the results observed on OpenLLM Leaderboard-v1 and v2.

Comparison with Leaderboard Results. The trends observed in MT-Bench do not always align with leaderboard outcomes. For instance, the Quantized Llama-3.1-405B model outperformed a newer Llama-3.3-70B-FP16 model on certain leaderboards, yet scored similarly or slightly lower in MT-Bench. Unlike the leaderboard tasks, which may not be particularly sensitive to *Coding* and *Math* challenges, the free-form and demanding nature of MT-Bench conversations highlights performance drops in more complex categories such as *Coding* and *STEM*. Thus, although computational metrics suggest that quantization does not uniformly degrade accuracy (RQ4), MT-Bench’s qualitative LLM-based evaluation reveals significant performance reductions in tasks known to be difficult.

RQ5 Findings: We observe that quantization considerably reduces performance in Coding and STEM tasks, while occasionally improving reasoning. The impact of quantization on multi-turn conversation quality does not consistently correlate with model size or type. Additionally, GPT4-based assessments sometimes misjudge incorrect answers as correct.

6 Conclusion

We evaluated instruction-tuned quantized LLMs across 13 datasets and 6 task types, using models ranging from 1B to 405B and 4 quantization methods, including GPTQ, AWQ, SmoothQuant, and FP8. We found that quantized LLMs generally outperformed smaller models in most tasks, except for hallucination detection and instruction-following. Performance varied by quantization method and precision, with weight-only quantization performing better in the 405B model. Task difficulty had little impact on accuracy loss. Our MT-Bench evaluation revealed that quantization significantly reduces performance in Coding and STEM tasks while occasionally enhancing reasoning. Additionally, GPT4-based assessments can misjudge incorrect answers as correct.

Acknowledgments

We thank the former Neubla ML team members — Minwook Ahn, Minhoo Park, Jinsol Kim, Raegeun Park, and Byonghwa Oh — for their valuable discussions and feedback. This work was supported by the Institute of Information & Communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT) (No.RS-2023-00277060, Development of open edge AI SoC hardware and software platform, No.RS-2024-00459797, Development of ML compiler framework for on-device AI, RS-2025-02214497, Development of low-level optimization program API technology for AI semiconductors)

Ethical Statement

There are no ethical issues.

References

- [Clark *et al.*, 2018] Peter Clark, Isaac Cowhey, Oren Etzioni, Tushar Khot, Ashish Sabharwal, Carissa Schoenick, and Oyvind Tafjord. Think you have solved question answering? try arc, the ai2 reasoning challenge. *arXiv preprint arXiv:1803.05457*, 2018.
- [Cobbe *et al.*, 2021] Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*, 2021.
- [contributors, 2024a] AutoAWQ contributors. Autoawq. <https://github.com/casper-hansen/AutoAWQ>, 2024.
- [contributors, 2024b] AutoGPTQ contributors. Autogptq. <https://github.com/AutoGPTQ/AutoGPTQ>, 2024.
- [Dettmers and Zettlemoyer, 2023] Tim Dettmers and Luke Zettlemoyer. The case for 4-bit precision: k-bit inference scaling laws. In *International Conference on Machine Learning*, pages 7750–7774. PMLR, 2023.
- [Dettmers *et al.*, 2022] Tim Dettmers, Mike Lewis, Younes Belkada, and Luke Zettlemoyer. Gpt3. int8 (): 8-bit matrix multiplication for transformers at scale. *Advances in Neural Information Processing Systems*, 35:30318–30332, 2022.
- [Dettmers *et al.*, 2024] Tim Dettmers, Artidoro Pagnoni, Ari Holtzman, and Luke Zettlemoyer. Qlora: Efficient finetuning of quantized llms. *Advances in Neural Information Processing Systems*, 36, 2024.
- [Dubey *et al.*, 2024] Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*, 2024.
- [Dutta *et al.*, 2024] Abhinav Dutta, Sanjeev Krishnan, Nipun Kwatra, and Ramachandran Ramjee. Accuracy is not all you need. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024.
- [Frantar *et al.*, 2022] Elias Frantar, Saleh Ashkboos, Torsten Hoefler, and Dan Alistarh. Optq: Accurate quantization for generative pre-trained transformers. In *The Eleventh International Conference on Learning Representations*, 2022.
- [Gong *et al.*, 2024] Ruihao Gong, Yang Yong, Shiqiao Gu, Yushi Huang, Chentao Lv, Yunchen Zhang, Xianglong Liu, and Dacheng Tao. Llmc: Benchmarking large language model quantization with a versatile compression toolkit, 2024.
- [Hendrycks *et al.*, 2021a] Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. Measuring massive multitask language understanding. *Proceedings of the International Conference on Learning Representations (ICLR)*, 2021.
- [Hendrycks *et al.*, 2021b] Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. Measuring mathematical problem solving with the math dataset. *arXiv preprint arXiv:2103.03874*, 2021.
- [Hong *et al.*, 2024] Junyuan Hong, Jinhao Duan, Chenhui Zhang, Zhangheng Li, Chulin Xie, Kelsey Lieberman, James Diffenderfer, Brian Bartoldson, Ajay Jaiswal, Kaidi Xu, et al. Decoding compressed trust: Scrutinizing the trustworthiness of efficient llms under compression. *arXiv preprint arXiv:2403.15447*, 2024.
- [Jaiswal *et al.*, 2023] Ajay Jaiswal, Zhe Gan, Xianzhi Du, Bowen Zhang, Zhangyang Wang, and Yinfei Yang. Compressing llms: The truth is rarely pure and never simple. *arXiv preprint arXiv:2310.01382*, 2023.
- [Jin *et al.*, 2024] Renren Jin, Jiangcun Du, Wuwei Huang, Wei Liu, Jian Luan, Bin Wang, and Deyi Xiong. A comprehensive evaluation of quantization strategies for large language models. *arXiv preprint arXiv:2402.16775*, 2024.
- [Kurtic *et al.*, 2024] Eldar Kurtic, Alexandre Marques, Shubhra Pandit, Mark Kurtz, and Dan Alistarh. ”give me bf16 or give me death”? accuracy-performance trade-offs in llm quantization. *arXiv preprint arXiv:2411.02355*, 2024.
- [Lewkowycz *et al.*, 2022] Aitor Lewkowycz, Anders Andreassen, David Dohan, Ethan Dyer, Henryk Michalewski, Vinay Ramasesh, Ambrose Slone, Cem Anil, Imanol Schlag, Theo Gutman-Solo, et al. Solving quantitative reasoning problems with language models. *Advances in Neural Information Processing Systems*, 35:3843–3857, 2022.
- [Li *et al.*, 2024] Shiyao Li, Xuefei Ning, Luning Wang, Tengxuan Liu, Xiangsheng Shi, Shengen Yan, Guohao Dai, Huazhong Yang, and Yu Wang. Evaluating quantized large language models. *arXiv preprint arXiv:2402.18158*, 2024.
- [Lin *et al.*, 2021] Stephanie Lin, Jacob Hilton, and Owain Evans. Truthfulqa: Measuring how models mimic human falsehoods. *arXiv preprint arXiv:2109.07958*, 2021.
- [Lin *et al.*, 2024] Ji Lin, Jiaming Tang, Haotian Tang, Shang Yang, Wei-Ming Chen, Wei-Chen Wang, Guangxuan Xiao, Xingyu Dang, Chuang Gan, and Song Han. Awq:

- Activation-aware weight quantization for on-device llm compression and acceleration. *Proceedings of Machine Learning and Systems*, 6:87–100, 2024.
- [Liu *et al.*, 2023] Peiyu Liu, Zikang Liu, Ze-Feng Gao, Dawei Gao, Wayne Xin Zhao, Yaliang Li, Bolin Ding, and Ji-Rong Wen. Do emergent abilities exist in quantized large language models: An empirical study. *arXiv preprint arXiv:2307.08072*, 2023.
- [Liu *et al.*, 2024] Yijun Liu, Yuan Meng, Fang Wu, Shenhao Peng, Hang Yao, Chaoyu Guan, Chen Tang, Xinzhu Ma, Zhi Wang, and Wenwu Zhu. Evaluating the generalization ability of quantized llms: Benchmark, analysis, and toolbox. *arXiv preprint arXiv:2406.12928*, 2024.
- [llmcompressor contributors, 2024] llmcompressor contributors. Llm compressor. <https://github.com/vllm-project/llm-compressor>, 2024.
- [Micikevicius *et al.*, 2022] Paulius Micikevicius, Dusan Stosic, Neil Burgess, Marius Cornea, Pradeep Dubey, Richard Grisenthwaite, Sangwon Ha, Alexander Heinecke, Patrick Judd, John Kamalu, et al. Fp8 formats for deep learning. *arXiv preprint arXiv:2209.05433*, 2022.
- [Namburi *et al.*, 2023] Satya Sai Srinath Namburi, Makesh Sreedhar, Srinath Srinivasan, and Frederic Sala. The cost of compression: Investigating the impact of compression on parametric knowledge in language models. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 5255–5273, 2023.
- [Rein *et al.*, 2023] David Rein, Betty Li Hou, Asa Cooper Stickland, Jackson Petty, Richard Yuanzhe Pang, Julien Dirani, Julian Michael, and Samuel R Bowman. Gpqa: A graduate-level google-proof q&a benchmark. *arXiv preprint arXiv:2311.12022*, 2023.
- [Sakaguchi *et al.*, 2021] Keisuke Sakaguchi, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. Winogrande: An adversarial winograd schema challenge at scale. *Communications of the ACM*, 64(9):99–106, 2021.
- [Sprague *et al.*, 2023] Zayne Sprague, Xi Ye, Kaj Bostrom, Swarat Chaudhuri, and Greg Durrett. Musr: Testing the limits of chain-of-thought with multistep soft reasoning. *arXiv preprint arXiv:2310.16049*, 2023.
- [Suzgun *et al.*, 2022] Mirac Suzgun, Nathan Scales, Nathanael Schärli, Sebastian Gehrmann, Yi Tay, Hyung Won Chung, Aakanksha Chowdhery, Quoc V Le, Ed H Chi, Denny Zhou, et al. Challenging big-bench tasks and whether chain-of-thought can solve them. *arXiv preprint arXiv:2210.09261*, 2022.
- [Team *et al.*, 2024] Gemma Team, Thomas Mesnard, Cassidy Hardin, Robert Dadashi, Surya Bhupatiraju, Shreya Pathak, Laurent Sifre, Morgane Rivière, Mihir Sanjay Kale, Juliette Love, et al. Gemma: Open models based on gemini research and technology. *arXiv preprint arXiv:2403.08295*, 2024.
- [Wan *et al.*, 2023] Zhongwei Wan, Xin Wang, Che Liu, Samiul Alam, Yu Zheng, Zhongnan Qu, Shen Yan, Yi Zhu, Quanlu Zhang, Mosharaf Chowdhury, et al. Efficient large language models: A survey. *arXiv preprint arXiv:2312.03863*, 1, 2023.
- [Wang *et al.*, 2024] Yubo Wang, Xueguang Ma, Ge Zhang, Yuansheng Ni, Abhranil Chandra, Shiguang Guo, Weiming Ren, Aaran Arulraj, Xuan He, Ziyang Jiang, et al. Mmlu-pro: A more robust and challenging multi-task language understanding benchmark. *arXiv preprint arXiv:2406.01574*, 2024.
- [Wei *et al.*, 2023] Xiuying Wei, Yunchen Zhang, Yuhang Li, Xiangguo Zhang, Ruihao Gong, Jinyang Guo, and Xianglong Liu. Outlier suppression+: Accurate quantization of large language models by equivalent and optimal shifting and scaling, 2023.
- [Xiao *et al.*, 2023] Guangxuan Xiao, Ji Lin, Mickael Seznec, Hao Wu, Julien Demouth, and Song Han. Smoothquant: Accurate and efficient post-training quantization for large language models. In *International Conference on Machine Learning*, pages 38087–38099. PMLR, 2023.
- [Xu *et al.*, 2024] Zhichao Xu, Ashim Gupta, Tao Li, Oliver Bentham, and Vivek Srikumar. Beyond perplexity: Multi-dimensional safety evaluation of llm compression. *arXiv preprint arXiv:2407.04965*, 2024.
- [Yao *et al.*, 2023] Zhewei Yao, Xiaoxia Wu, Cheng Li, Stephen Youn, and Yuxiong He. Zeroquant-v2: Exploring post-training quantization in llms from comprehensive study to low rank compensation. *arXiv preprint arXiv:2303.08302*, 2023.
- [Zellers *et al.*, 2019] Rowan Zellers, Ari Holtzman, Yonatan Bisk, Ali Farhadi, and Yejin Choi. Hellaswag: Can a machine really finish your sentence? *arXiv preprint arXiv:1905.07830*, 2019.
- [Zhang *et al.*, 2024] Yue Zhang, Ming Zhang, Haipeng Yuan, Shichun Liu, Yongyao Shi, Tao Gui, Qi Zhang, and Xuanjing Huang. Llmeval: A preliminary study on how to evaluate large language models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 19615–19622, 2024.
- [Zheng *et al.*, 2023] Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, et al. Judging llm-as-a-judge with mt-bench and chatbot arena. *Advances in Neural Information Processing Systems*, 36:46595–46623, 2023.
- [Zhou *et al.*, 2023] Jeffrey Zhou, Tianjian Lu, Swaroop Mishra, Siddhartha Brahma, Sujoy Basu, Yi Luan, Denny Zhou, and Le Hou. Instruction-following evaluation for large language models, 2023.
- [Zhu *et al.*, 2023] Xunyu Zhu, Jian Li, Yong Liu, Can Ma, and Weiping Wang. A survey on model compression for large language models. *arXiv preprint arXiv:2308.07633*, 2023.

Appendices

Table of Contents

A	Details of Datasets	11
B	Additional Models	12
C	Detailed Results of MT-Bench	12
C.1	Category-Level Analysis	12
C.2	Model-Level Analysis and Multi-Turn Results	12
C.3	Cost for MT-Bench	13
C.4	Qualitative Analysis	13
C.5	Misjudged Cases	14
C.6	A Guide to Effective MT-Bench Usage	15
D	Additional Details for Reproducibility	15
D.1	Experimental Configuration	15
D.2	Details on how leaderboard scores are calculated	15
D.3	Quantization Method	16
D.4	Quantization Tools and Configurations	16
D.5	Calibration Data for Quantization	17

List of Figures

3	GPTQ Configuration	17
4	AutoAWQ Configuration	17
5	LLM Compressor’s GPTQ Configuration	17
6	LLM Compressor’s SmoothQuant Configuration	18
7	LLM Compressor’s FP8 Configuration	18

List of Tables

3	Overview of different datasets, their type, # of tasks, # of samples, evaluation ability, and assessment metrics. ICL, CoT, and IF are acronyms for In-Context Learning, Chain-of-Thought, and Instruction Following, respectively, which are emergent abilities.	11
4	Evaluation of Vicuna and Gemma on OpenLLM Leaderboard-v1 tasks including ARC-c, GSM8k, HellaSwag, MMLU, TruthfulQA, and Winogrande. * denotes the use of AutoGPTQ for quantization.	12
5	Evaluation of Vicuna and Gemma on OpenLLM Leaderboard-v2 tasks including BBH, GPQA, IFEval, Math-Lvl-5, MMLU-PRO, and MuSR. * denotes the use of AutoGPTQ for quantization.	12
6	For each category, the MT-Bench scores (FP16, GPTQ, AWQ) are presented for 12 models. The difference in MT-Bench scores compared to FP16 is computed and shown for GPTQ and AWQ. The three categories with the largest score drops compared to FP16 are highlighted in red, while The categories with score increases are highlighted in green.	19
7	A evaluation of the influence of quantization on the MT-Bench multi-turn conversation benchmark. * denotes the use of AutoGPTQ for quantization. The API cost represents the GPT4 usage fee incurred when running MT-Bench for a single model. We highlighted the best value in green.	20
8	MT-Bench (single mode) results for pretrained models without intruction tuning.	20
9	Calibration settings for quantization, including dataset used, sample count, fixed sequence length, and default configurations ensure stable performance across models. * denotes the use of AutoGPTQ for quantization, while ** uses vLLM project’s LLM-Compressor for GPTQ quantization. Also, *** denotes the use of Intel’s AutoRound for GPTQ quantization.	21

A Details of Datasets

The 12 datasets used in the experiments are summarized in Table 3. Each dataset is described in detail below.

Common-sense Reasoning. Among these, ARC, HellaSwag, and Winogrande target common-sense reasoning that is evident to humans yet remains challenging for AI. Notably, for ARC, we evaluated only the *challenge* subset to focus on the most challenging questions.

MMLU. The MMLU dataset provides multiple-choice questions spanning 57 subjects, ranging from STEM fields to the humanities. Each question has four answer choices, and the dataset tests a model’s ability to handle a wide variety of topics with moderate to high difficulty.

GPQA. GPQA consists of PhD-level queries formulated by domain experts. Since many questions require specialized knowledge, this dataset is particularly challenging for models lacking deep understanding in advanced subjects.

MMLU-PRO. MMLU-PRO is an extended variant of MMLU. It expands each question from four to ten answer choices and covers 14 topics, thereby increasing difficulty and necessitating more nuanced reasoning.

BBH. BBH comprises 23 subjects, including multi-step arithmetic, algorithmic reasoning, and satire. It evaluates a model’s capacity for human-like language understanding and logical processing.

MuSR. MuSR is divided into three tasks: murder mysteries, object placements, and team allocation. Solving these problems relies heavily on Chain-of-Thought (CoT) capabilities, as models must parse lengthy contexts, reason step by step, and combine partial conclusions.

IFEval. IFEval tests whether models accurately follow instructions. It focuses purely on adherence to directives rather than content generation. For example, if a prompt requires writing over 400 words, the model should satisfy this requirement regardless of what it writes.

TruthfulQA. TruthfulQA evaluates how factually correct a model’s response is, serving as a benchmark for gauging hallucination tendencies. Models that produce incorrect or misleading outputs are penalized in this dataset.

GSM8K and MATH-Lvl-5. GSM8K includes elementary-level math problems that generally need multi-step reasoning. MATH-Lvl-5 contains high-school competition problems across seven subjects, requiring LaTeX notation for equations and Asymptote for figures. These problems test a model’s ability to handle complex calculations and formatting.

MT-Bench. MT-Bench employs GPT4 as a judge, using a single-answer grading system to evaluate multi-turn dialogues across eight categories (e.g., Writing, Roleplay, Reasoning). This free-form conversation benchmark emphasizes coherence, depth, and adaptability, making it well-suited for assessing real-world usage scenarios.

Alignment with OpenLLM Leaderboards. We selected 12 benchmarks that align with the Huggingface OpenLLM Leaderboard-v1 (covering April 2023–June 2024) and Leaderboard-v2 (launched on June 26, 2024). This alignment helps reduce data contamination in newer models. For each task, we adopted the same few-shot In-Context Learning (ICL) settings used in the leaderboards [Liu *et al.*, 2023; Li *et al.*, 2024], ensuring consistent comparisons. Since Leaderboard-v2 does not yet include quantized model results and Leaderboard-v1 only features limited GPTQ examples, we performed all measurements in the same environment [Gong *et al.*, 2024] for a fair comparison.

Overall, these datasets cover a wide spectrum of reasoning abilities, knowledge domains, and linguistic complexities, allowing us to rigorously evaluate how quantization impacts instruction-tuned LLMs across different tasks.

Type	Dataset	# Tasks	# Test Samples	Evaluation Ability	Metrics
CommonSenseQA	ARC-Challenge [Clark <i>et al.</i> , 2018]	1	1,172	Primary School Knowledge, ICL	Accuracy ↑
	HellaSwag [Zellers <i>et al.</i> , 2019]	1	10,003	Knowledge, ICL	Accuracy ↑
	Winogrande [Sakaguchi <i>et al.</i> , 2021]	1	1,767	Knowledge, ICL	Accuracy ↑
Complex Knowledge & Language Understanding	MMLU [Hendrycks <i>et al.</i> , 2021a]	57	14,042	Knowledge, ICL	Accuracy ↑
	GPQA [Rein <i>et al.</i> , 2023]	3	448	PhD-Lvl Knowledge	Accuracy ↑
	MMLU-PRO [Wang <i>et al.</i> , 2024]	13	12,032	Expert-Lvl Knowledge, ICL	Accuracy ↑
	BBH [Suzgun <i>et al.</i> , 2022]	23	6,511	Knowledge, Algorithmic Reasoning, Language Understanding, ICL	Accuracy ↑
	MuSR [Sprague <i>et al.</i> , 2023]	3	756	Reasoning, CoT	Accuracy ↑
Instruction Following	IFEval [Zhou <i>et al.</i> , 2023]	1	541	Rigorousness, IF	Strict Accuracy ↑
Hallucination	TruthfulQA [Lin <i>et al.</i> , 2021]	38	817	Truthfulness	Multi-True (MC2) ↑
Mathematics	GSM8K [Cobbe <i>et al.</i> , 2021]	1	1,319	Mathematical Reasoning, CoT	Strict-Match ↑
	MATH-Lvl-5 [Hendrycks <i>et al.</i> , 2021b; Lewkowycz <i>et al.</i> , 2022]	7	12,000	Mathematical Reasoning, CoT	Exact-Match ↑
Dialogue	MT-Bench [Zheng <i>et al.</i> , 2023]	8	80	Conversation Quality	Score (0–10) ↑

Table 3: Overview of different datasets, their type, # of tasks, # of samples, evaluation ability, and assessment metrics. ICL, CoT, and IF are acronyms for In-Context Learning, Chain-of-Thought, and Instruction Following, respectively, which are emergent abilities.

B Additional Models

Table 4 and Table 5 demonstrate how GPTQ and AWQ quantization methods affect the accuracy of Vicuna and Gemma across OpenLLM Leaderboard-v1 and OpenLLM Leaderboard-v2. Overall, AWQ tends to incur smaller accuracy drops than GPTQ when compared to the FP16 baseline, particularly on tasks such as ARC-c and GSM8k. In contrast, GPTQ often exhibits more noticeable performance declines. This pattern persists in both Leaderboard-v1 (covering ARC-c, GSM8k, HellaSwag, MMLU, TruthfulQA, Winogrande) and Leaderboard-v2 (covering BBH, GPQA, IFEval, Math-Lvl-5, MMLU-PRO, MuSR), suggesting that AWQ maintains more stable results across a variety of tasks.

OpenLLM Leaderboard-v1 ↑										
Model	Method	W/A	Storage (GB)	ARC-c (25-shot) acc_norm	GSM8k (5-shot) acc	HellaSwag (10-shot) acc_norm	MMLU (5-shot) acc	TruthfulQA (0-shot) mc2	Winogrande (5-shot) acc	Avg.
Vicuna-7B-v1.3	FP16	16 / 16	14	51.28	13.34	77.36	47.21	47.01	71.03	51.21
	GPTQ*	4 / 16	3.5	46.84 (↓4.44)	10.31 (↓3.03)	75.06 (↓2.30)	43.54 (↓3.67)	44.18 (↓2.83)	69.06 (↓1.97)	48.17 (↓3.04)
	AWQ	4 / 16	3.5	51.20 (↓0.08)	12.74 (↓0.60)	76.63 (↓0.73)	46.34 (↓0.87)	46.65 (↓0.36)	71.35 (↑0.32)	50.82 (↓0.39)
Gemma-2B-it	FP16	16 / 16	4	44.03	4.01	63.01	36.92	45.76	61.48	42.54
	GPTQ*	4 / 16	1	44.28 (↑0.25)	3.10 (↓0.91)	61.76 (↓1.25)	37.21 (↑0.29)	46.31 (↑0.55)	60.14 (↓1.34)	42.13 (↓0.41)
	AWQ	4 / 16	1	42.74 (↓1.29)	1.66 (↓2.35)	60.98 (↓2.03)	36.21 (↓0.71)	45.18 (↓0.58)	59.66 (↓1.82)	41.07 (↓1.47)
Gemma-7B-it	FP16	16 / 16	14	51.54	27.45	72.06	51.62	47.05	64.72	52.41
	GPTQ*	4 / 16	3.5	51.02 (↓0.52)	25.62 (↓1.83)	70.75 (↓1.31)	51.21 (↓0.41)	47.11 (↑0.06)	66.29 (↑1.57)	52.00 (↓0.41)
	AWQ	4 / 16	3.5	50.43 (↓1.11)	25.85 (↓1.60)	71.27 (↓0.79)	51.14 (↓0.48)	46.11 (↓0.94)	66.37 (↑1.65)	51.86 (↓0.55)

Table 4: Evaluation of Vicuna and Gemma on OpenLLM Leaderboard-v1 tasks including ARC-c, GSM8k, HellaSwag, MMLU, TruthfulQA, and Winogrande. * denotes the use of AutoGPTQ for quantization.

OpenLLM Leaderboard-v2 ↑										
Model	Method	W/A	Storage (GB)	BBH (3-shot)	GPQA (0-shot)	IFEval (0-shot)	Math-Lvl-5 (4-shot)	MMLU-PRO (5-shot)	MuSR (0-shot)	Avg.
Vicuna-7B-v1.3	FP16	16 / 16	14	12.26	2.08	26.13	1.61	10.03	5.72	9.64
	GPTQ*	4 / 16	3.5	7.61 (↓4.65)	0.12 (↓1.96)	28.07 (↑1.94)	0.79 (↓0.82)	6.40 (↓3.63)	3.85 (↓1.87)	7.81 (↓1.83)
	AWQ	4 / 16	3.5	10.65 (↓1.61)	0.61 (↓1.47)	22.90 (↓3.23)	0.82 (↓0.79)	8.52 (↓1.51)	3.15 (↓2.57)	7.78 (↓1.86)
Gemma-2B-it	FP16	16 / 16	4	4.67	6.72	27.02	2.01	6.23	4.82	8.58
	GPTQ*	4 / 16	1	4.16 (↓0.51)	6.72 (0.00)	28.76 (↑1.74)	1.38 (↓0.63)	5.94 (↓0.29)	4.81 (↓0.01)	8.63 (↓0.05)
	AWQ	4 / 16	1	4.96 (↑0.29)	3.79 (↓2.93)	29.02 (↑2.00)	1.39 (↓0.62)	5.69 (↓0.54)	3.92 (↓0.90)	8.13 (↓0.45)
Gemma-7B-it	FP16	16 / 16	14	15.89	3.54	35.68	2.65	14.55	17.00	14.89
	GPTQ*	4 / 16	3.5	13.51 (↓2.38)	9.65 (↑6.11)	35.56 (↓0.12)	1.92 (↓0.73)	13.66 (↓0.89)	18.26 (↑1.26)	15.43 (↑0.54)
	AWQ	4 / 16	3.5	15.63 (↓0.26)	6.23 (↑2.69)	36.39 (↑0.71)	2.40 (↓0.25)	13.76 (↓0.79)	15.73 (↓1.27)	15.02 (↑0.13)

Table 5: Evaluation of Vicuna and Gemma on OpenLLM Leaderboard-v2 tasks including BBH, GPQA, IFEval, Math-Lvl-5, MMLU-PRO, and MuSR. * denotes the use of AutoGPTQ for quantization.

C Detailed Results of MT-Bench

C.1 Category-Level Analysis

Table 6 shows that the Reasoning category improves for 9 out of 12 models after quantization. By contrast, the categories with the largest accuracy drops are Coding, Humanities, and Extraction. Among the 24 total quantized configurations (Models × Quantization), Coding, Humanities, and Extraction each appear 16, 12, and 10 times, respectively, among the top three categories with the most severe performance losses. When analyzing all eight MT-Bench categories under quantization, the most pronounced declines are observed in Math, Coding, and Reasoning. Llama-3.3-70B, the newest model, achieves notably higher scores across all eight categories.

C.2 Model-Level Analysis and Multi-Turn Results

Table 7 presents the multi-turn performance for each model. While smaller models such as Llama-3.1-8B-it, Llama-3.2-1B, and Llama-3.2-3B sometimes show unexpected boosts in Reasoning (e.g., +1.65 in Llama-3.2-3B), larger models generally do not experience such gains. Although Llama-3.2 is among the latest releases, it does not exhibit markedly different multi-turn behavior compared to other recent models. Additionally, models like Llama-3.1 tend to rank highly in MT-Bench, possibly because of their similarity to GPT4 in certain tasks; however, these close scores make it difficult to treat the results as definitive. Quantized variants also perform well under this evaluation scheme.

C.3 Cost for MT-Bench

Across all models, the MT-Bench evaluation incurs a cost ranging from \$4 to \$7 each, totaling about \$200 in GPT4 API usage. This highlights the resource-intensive nature of multi-turn LLM evaluations, particularly when using state-of-the-art judging models.

C.4 Qualitative Analysis

An improved example and a degraded example of the Reasoning question on MT-Bench by the quantized Llama-3.2-3B-it.

QUESTION:

Turn 1: You can see a beautiful red house to your left and a hypnotic greenhouse to your right, an attractive heated pink place in the front. So, where is the White House?

Turn 2: Does the original question contain any clues to definitively determine the location of the White House?

reference:

Turn 1: The answer is Washington, DC.

Turn 2: No.

FP16 (original) ANSWER:

Turn 1 (correct): The White House is actually located in [Washington D.C.](#), which is the capital city of the United States. [...] The red house you are referring to is likely the White House itself, ...

→ **GPT4 Judgement: (6/10)** ... If the references to the red house, greenhouse, and heated pink place are part of a riddle or a creative description, **additional context would be helpful** to provide a more accurate interpretation.

Turn 2 (generally correct):

... So, **while the original question does not contain any explicit clues**, I think it is possible that the beautiful red house, hypnotic greenhouse, ...

→ **GPT4 Judgement: (2/10)** ... **Descriptions such as a beautiful red house, a hypnotic greenhouse, and an attractive heated pink place do not relate to its actual location.** ...

W4A16-GPTQ ANSWER (Degraded):

Turn 1 (correct): [The White House is actually located in Washington, D.C.](#), [...] {But if you're looking for a building that is similar in style and architecture to the White House, you might be thinking of the U.S. Capitol Building's neighbor, the Jefferson Memorial's neighbor, the Arlington National Cemetery's neighbor... } × 11 (This part is repeated 11 times).

→ **GPT4 Judgement: (1/10)** ... **the assistant provided a long, repetitive, and confusing response that did not directly answer the question.** ...

Turn 2 (wrong): ... **these hints suggest the White House is in Washington, D.C.** While these interpretations are educated guesses, they align with the clues provided.

→ **GPT4 Judgement: (1/10)** ... The assistant incorrectly interprets the descriptions of a "beautiful red house", a "hypnotic greenhouse", and an "attractive heated pink place" as clues to the location of the White House. **These descriptions are unrelated to the White House's location.**

W4A16-AWQ ANSWER (Improved):

Turn 1 (correct): [The White House is actually located in Washington, D.C.](#), ...

→ **GPT4 Judgement: (10/10)** ... It provides the same information as the reference answer, ...

Turn 2 (correct): ... **these descriptions don't provide any specific clues to determine the location of the White House.** ...

→ **GPT4 Judgement: (10/10)** ... It accurately states that the original question does not contain any clues to definitively determine the location of the White House. ...

Manual inspection: The original model correctly answered turn 1 but lost four points due to excessive supplementary explanations and unnecessary use of technical terms. Similarly, turn 2 was correct but earned only two points for the same reason. GPTQ provided the correct answer for turn 1 but was penalized for being overly lengthy and repetitive, and it failed to produce the correct answer for turn 2. In contrast, AWQ delivered concise and clear and concise responses, achieving a perfect score.

C.5 Misjudged Cases

A misjudged example of the Reasoning question on MT-Bench by the quantized Llama-3.2-3B-it.

QUESTION:

Turn 1: David has three sisters. Each of them has one brother. How many brothers does David have?

Turn 2: If we change the previous question and assume that each sister of David has two brothers, how many brothers would David have?

reference:

Turn 1: David has no brother. He is the one brother of his three sisters.

Turn 2: David has one brother.

FP16 (original) ANSWER:

Turn 1 (wrong): That's a clever question! Since each of David's sisters has one brother, and David has three sisters, that means David has three brothers. The answer is three.

→ **GPT4 Judgement: (2/10) The assistant's answer is incorrect.** The question states that each of David's sisters has one brother. This implies that David is the only brother to his three sisters. Therefore, David does not have any brothers. The assistant incorrectly interpreted the question and provided an incorrect answer. The correct answer should be that David has no brothers.

Turn 2 (wrong): Let's analyze the situation again. If each of David's sisters has two brothers, and David has three sisters, that means each sister has two brothers. However, David is also one of those brothers. So, when we count the brothers, we need to include David as well. Therefore, David has three brothers in total.

→ **GPT4 Judgement: (2/10) The assistant's response to the second question is incorrect.** The assistant incorrectly states that David has three brothers when the correct answer is that David has two brothers. The assistant seems to be confused about the concept of siblings and is counting David as one of his own brothers, which is incorrect. The correct response should be: If each of David's sisters has two brothers, then David would have two brothers.

W4A16-AWQ ANSWER:

Turn 1 (wrong): Let's break it down step by step: 1. David has three sisters. 2. Each of his sisters has one brother. 3. Since each sister has one brother, David must have one brother as well. So, David has one brother.

→ **GPT4 Judgement: (10/10) The assistant's answer is correct and aligns with the reference answer.** The assistant provides a detailed explanation, breaking down the problem step by step, which can be helpful for the user to understand the logic behind the answer. There are no mistakes to correct.

Turn 2 (wrong): Let's re-examine the situation: 1. David has three sisters. 2. Each of his sisters has two brothers. 3. Since each sister has two brothers, David must have two brothers as well. So, David has two brothers.

→ **GPT4 Judgement: (10/10) The assistant's answer is correct.** The assistant correctly identified that if each of David's sisters has two brothers, then David would have two brothers. The assistant's explanation is clear and easy to understand, breaking down the situation step by step. There are no mistakes to correct in the assistant's response.

Manual inspection: Despite being provided with a reference answer (he (David) is the one brother), GPT4 appears to infer and evaluate a new answer independently. This question's correct answer varies depending on David's gender: if male, the answer matches the reference, while if female, the correct answers are "one" and "two," respectively. GPT4 seems to interpret David's gender differently with each evaluation. Nevertheless, in the case of FP16, the model produced answers significantly distant from the correct ones, whereas the quantized model provided an answer that could be correct under the assumption that David is female.

A misjudged example of the STEM question on MT-Bench by the quantized Llama-3.2-3B-it.

QUESTION: Consider a satellite that is in a circular orbit around the Earth. The speed of the satellite decreases. What will happen to the satellite’s orbital radius and period of revolution? Please justify your answer using principles of physics.
reference: The orbital radius will increase and the period of revolution will increase.

FP16 (original) ANSWER:(correct) ... In summary, if the speed of a satellite in a circular orbit around the Earth decreases, its orbital radius will increase, and its period of revolution will decrease. ...

→ **GPT4 Judgement: (1/10)** ... Consequently, the orbital radius will decrease, not increase. [...] If the orbital radius decreases, the satellite will complete its orbit in less time, meaning the period of revolution will also decrease. ...

W4A16-GPTQ ANSWER(correct): ... To maintain its orbit, the satellite needs to increase its distance from the Earth, implying an increase in orbital radius. [...] Consequently, the period of revolution increases. ...

→ **GPT4 Judgement: (2/10)** ... When a satellite’s speed decreases, [...] it will fall into a lower orbit, which means its orbital radius will decrease. [...] As the satellite’s speed decreases and it falls into a lower orbit, the time it takes to complete one orbit (the period of revolution) will also decrease.

Manual inspection: Despite the reference answer being provided, GPT4 guided incorrect answers through its own reasoning and justification, ultimately misjudging the responses of FP16 and GPTQ models.

C.6 A Guide to Effective MT-Bench Usage

The dialogue evaluation using MT-Bench is useful when utilizing models designed for chat. As listed in the table 8, when evaluating the quality of free-form text generation, models that are not instruction-tuned for chat tend to produce lower evaluation scores. Therefore, to conduct a meaningful evaluation, MT-Bench should be used with LLMs that have been instruction-tuned for chat.

D Additional Details for Reproducibility

D.1 Experimental Configuration

Random Seed Configuration

The lm-eval random seed was set to the default. Consequently, the Python seed, NumPy seed, Torch seed, and few-shot seed were set to 0, 1234, 1234, and 1234, respectively.

Decoding Settings

In lm-eval, if no specific settings are provided, decoding defaults to greedy decoding, which we used. For vllm, temperature, top_k, and top_p were configured to operate in *greedy decoding* mode, using the default settings.

Package Versions

Accuracy may vary depending on the versions of the packages used. Below are the versions we used in our experiments:

- Accelerate: 1.2.1
- Transformers: 4.47.1
- Vllm: 0.6.6.post1
- Llmcompressor: 0.3.1
- Auto-gptq: 0.7.1
- Autoawq: 0.2.7.post3
- Bitsandbytes: 0.45.0
- lm_eval: 0.4.4
- fschat (mt-bench): 0.2.36

D.2 Details on how leaderboard scores are calculated

OpenLLM Leaderboard-v1

To compute scores on OpenLLM Leaderboard v1, we followed guidelines described in leaderboard-v1 website¹. We evaluated models on six key benchmarks using the Eleuther AI Language Model Evaluation Harness, a unified framework designed to assess generative language models across diverse tasks. The specific metric and few-shot parameters used are as follows:

¹https://huggingface.co/docs/leaderboards/open_llm_leaderboard/archive

- ARC: 25-shot, arc-challenge (acc_norm)
- HellaSwag: 10-shot, hellaswag (acc_norm)
- TruthfulQA: 0-shot, truthfulqa (mc2)
- MMLU: 5-shot (acc)
- Winogrande: 5-shot, winogrande (acc)
- GSM8k: 5-shot, gsm8k (acc)

OpenLLM Leaderboard-v2

To compute scores on OpenLLM Leaderboard v2, we followed guidelines described in leaderboard-v2 website ².

Our evaluation protocol spans multiple datasets and metrics. We measure 0-shot instance- and prompt-level strict accuracy for **IFEval**. **Big Bench Hard (BBH)** is assessed in a 3-shot setting, using normalized accuracy across various subtasks with different `num_choices`. **Math Challenges** (Math Level 5) employs a 4-shot exact-match metric. **GPQA** is tested under 0-shot conditions with a normalized accuracy measure and four answer choices. **MUSR** includes sub-tasks such as Murder Mysteries, Object Placement, and Team Allocation, each evaluated with 0-shot normalized accuracy. Lastly, **MMLU-PRO** is examined with a 5-shot approach, measuring accuracy out of ten multiple-choice options. By adhering to these configurations, researchers can reproduce our reported results and further advance quantization research.

We describe here the methodology for normalizing scores on the OpenLLM Leaderboard v2, which also leverages the `lm_eval` framework for evaluation. This normalization procedure aims to account for varying task difficulties and random-guess baselines, thereby providing a consistent (0–100) scale that supports meaningful comparisons across tasks.

General Normalization Process. Scores are normalized in two steps: subtracting a lower bound (i.e., the random baseline score) and then scaling the outcome to a range of 0–100 using:

$$\text{normalized_score} = 100 \times \frac{(\text{raw_score} - \text{lower_bound})}{(\text{higher_bound} - \text{lower_bound})}.$$

For tasks without subtasks (e.g., GPQA, MMLU-PRO), the lower bound corresponds to the inverse of `num_choices`, and the higher bound is set to 1. Tasks with multiple subtasks (e.g., MUSR, BBH) require per-subtask normalization prior to averaging the resulting scores.

Generative Evaluations. For generative tasks such as MATH and IFEval, we adopt an exact-match or strict-accuracy approach, where the lower bound is effectively 0, reflecting the improbability of generating correct answers at random. Specifically, MATH uses exact-match accuracy, and IFEval measures strict accuracy at both instance- and prompt-level evaluations.

This normalization code could be referred to the Colab ³.

D.3 Quantization Method

To evaluate the performance of quantized models, we applied various quantization methods, including GPTQ [Frantar *et al.*, 2022], AWQ [Lin *et al.*, 2024], SmoothQuant [Xiao *et al.*, 2023], and the FP8 [Micikevicius *et al.*, 2022]. The considered quantization methods fall under the category of Post-Training Quantization (PTQ), with GPTQ and AWQ being weight-only quantization techniques. GPTQ employed layer-wise quantization and leverages inverse Hessian information to update weights, mitigating accuracy loss. We used both AutoGPTQ [contributors, 2024b] and llmcompressor tools to apply GPTQ to LLMs, as both tools offer support for this method. This allowed us to further analyze the accuracy differences resulting from the implementation of the GPTQ algorithm in each tool. The group size used for GPTQ quantization in both tools is 128.

Activation-Aware Weight Quantization (AWQ) was designed to effectively quantize large language models (LLMs) while preserving the precision of the most critical weights. We used AutoAWQ [contributors, 2024a], an extended implementation of AWQ designed to make it easier to apply, and the configured group size is 128.

SmoothQuant proposed a per-channel scaling approach, shifting the complexity of high-precision quantization from activations to weights. This allowed for 8-bit quantization of both weights and activations in LLMs. In the case of SmoothQuant, we applied it to LLMs using the llmcompressor [llmcompressor contributors, 2024]. For FP8 quantization, we also used llmcompressor, implementing it with the E4M3 format. The FP8 E4M3 format is directly supported by NVIDIA’s Hopper and Ada Lovelace architectures and is compatible with the vLLM library. More details about the quantization configurations are in Appendix D.4.

D.4 Quantization Tools and Configurations

The considered quantization methods in this study include GPTQ, AWQ, SmoothQuant, and FP8, with the tools used being AutoGPTQ, AutoAWQ, and llmcompressor. Each tool and method follows specific configurations to ensure optimal performance.

²https://huggingface.co/docs/leaderboards/open_llm_leaderboard/normalization

³https://colab.research.google.com/drive/1-aPrFJwdifhVLxzJcsYXeebqNi_5vaw?usp=sharing

AutoGPTQ

For AutoGPTQ, the model is quantized to 4-bit with a group size of 128, which is the recommended value. Symmetric quantization is applied (`sym=True`), allowing zero to be precisely represented, which can offer speedups. The activation descent is set to `True` (`desc_act=True`), but setting it to `False` can significantly speed up inference at the cost of slightly reduced perplexity. A 10% damping factor is used (`damp_percent=0.1`) to further refine the performance. Figure 3 shows a configuration for GPTQ quantization.

```
1 quantize_config =  
2 BaseQuantizeConfig(bits=4, group_size=128, desc_act=True, sym=True, damp_percent=0.1)
```

Figure 3: GPTQ Configuration

AutoAWQ

In AutoAWQ, asymmetric quantization scheme is employed with a group size of 128, and the model is quantized to 4-bit. This configuration is optimized for maintaining accuracy while reducing computational complexity. Figure 4 shows a configuration for AWQ quantization.

```
1 quant_config={"zero_point": True, "q_group_size": 128, "w_bit": 4, "version": "GEMM"}
```

Figure 4: AutoAWQ Configuration

vLLM’s llmcompressor

For llmcompressor with GPTQ, only the weights of the linear operators within the transformer blocks are quantized. Symmetric per-channel quantization is applied, where a linear scaling per output dimension maps the INT8 or INT4 representations to floating-point weights. AutoGPTQ is also used for this process, with a 10% damping factor for further precision.

In llmcompressor with SmoothQuant, symmetric per-channel quantization is again applied, focusing on the weights of the linear operators. This method uses INT8 quantization, and the activations are also quantized using INT8, ensuring consistency across both weights and activations.

Lastly, llmcompressor with FP8 utilizes FP8 types, which have two distinct representations typically supported by hardware. FP8 (E4M3) consists of 1 sign bit, 4 exponent bits, and 3 bits of mantissa, capable of storing values up to +/-448. Both weights and activations are quantized per tensor using symmetric quantization, ensuring uniform precision across the model.

Figure 5, Figure 6, and Figure 7 show the configurations for GPTQ, SmoothQuant, and FP8, respectively.

```
1 recipe = GPTQModifier(targets="Linear", ignore=["lm_head"], dampening_frac=0.01,  
2 scheme="W4A16" or "W8A16")
```

Figure 5: LLM Compressor’s GPTQ Configuration

D.5 Calibration Data for Quantization

To ensure the reproducibility of our experiments, Table 9 details the calibration settings used for quantization.

We employed different quantization tools for each model, as tool compatibility varies across models. Quantization was carried out using the default settings of each respective tool to maintain consistent performance. While these default configurations were effective for most tasks, specific parameters such as group size and dataset selection were adjusted based on the quantization algorithm and the model architecture. As discussed in prior work [Liu *et al.*, 2024], the accuracy of quantized models can be influenced by the variety of subjects included in the calibration data. However, a thorough investigation of calibration datasets is beyond the scope of our study.

Below, we provide descriptions of each dataset used in our evaluations.

- **Wikitext2**⁴: It is a default dataset used in AutoGPTQ tool.
- **mit-han-lab/pile-val-backup**⁵: In the AWQ paper, the default calibration dataset (autoAWQ) is a subset of the Pile.

⁴<https://huggingface.co/datasets/mindchain/wikitext2>

⁵<https://huggingface.co/datasets/mit-han-lab/pile-val-backup>

```

1 # default scheme
2 recipe = [SmoothQuantModifier(smoothing_strength=0.8),
3 GPTQModifier(scheme="W8A8", targets="Linear", ignore=["lm_head"]),
4 GPTQModifier(sequential=True, targets="Linear",
5   scheme="W8A8", ignore=["lm_head"],
6   dampening_frac=0.01, observer="mse")]
7 # for Llama-3.2
8 recipe = [SmoothQuantModifier(smoothing_strength=0.7,
9   mappings=[[["re:.*q_proj", "re:.*k_proj", "re:.*v_proj"],
10    "re:.*input_layernorm"], [{"re:.*gate_proj", "re:.*up_proj"},
11    "re:.*post_attention_layernorm"], [{"re:.*down_proj"}, "re:.*up_proj"], ],),
12 GPTQModifier(sequential=True, targets="Linear", scheme="W8A8",
13   ignore=["lm_head"], dampening_frac=0.01)]

```

Figure 6: LLM Compressor’s SmoothQuant Configuration

```

1 # default scheme
2 recipe = "" quant_stage: quant_modifiers: QuantizationModifier: ignore: ["lm_head"]
3   config_groups: group_0: weights: num_bits: 8 type: float strategy: tensor
4   dynamic: false symmetric: true input_activations: num_bits: 8 type: float
5   strategy: tensor dynamic: false symmetric: true targets: ["Linear"]""
6 #for Llama-3.2
7 recipe = QuantizationModifier(targets="Linear", scheme="FP8", ignore=["lm_head"])

```

Figure 7: LLM Compressor’s FP8 Configuration

- **HuggingFaceH4/ultrachat-200k**⁶: UltraChat 200k is a heavily filtered subset of the original UltraChat dataset, comprising 200k dialogues.
- **LLM-compression-calibration**⁷: The dataset was constructed by extracting 10k representative samples from OpenPlatypus for use in quantization.
- **Random data from vocab of tokenizer**: The calibration dataset was generated by randomly sampling 256 sequences, each containing 8,192 tokens uniformly drawn from the model’s vocabulary.
- **NeelNanda/pile-10k**⁸: The AutoRound tool uses a subset of the Pile as its default dataset.

⁶https://huggingface.co/datasets/HuggingFaceH4/ultrachat_200k

⁷https://huggingface.co/datasets/neuralmagic/LLM_compression_calibration

⁸<https://huggingface.co/datasets/NeelNanda/pile-10k>

Model	Method	Writing	Roleplay	Reasoning	Math	Coding	Extraction	STEM	Humanities	Avg.
Vicuna-7B-v1.3	FP16	8.10	7.45	4.65	2.30	3.55	5.00	7.83	9.10	6.00
	GPTQ	-0.05	-0.50	+0.20	-0.40	-0.95	+0.45	-0.53	-0.55	-0.29
	AWQ	+0.20	+0.30	+0.70	+0.75	-1.00	-0.55	+0.07	-0.20	+0.03
Gemma-2B-it	FP16	6.67	5.90	2.75	1.75	2.40	4.45	6.05	7.10	4.63
	GPTQ	-0.42	-0.20	+0.30	+0.55	-0.65	+0.15	-0.25	-0.32	-0.10
	AWQ	-0.70	+0.45	-0.05	+0.05	-0.75	-0.65	+0.30	-0.25	-0.20
Gemma-7B-it	FP16	6.97	6.60	5.00	3.70	3.75	5.90	7.05	8.93	5.99
	GPTQ	+0.11	-0.60	-0.15	-1.10	-0.70	+0.35	-0.13	-1.18	-0.43
	AWQ	-0.09	-0.70	-0.35	-0.90	-0.40	-0.40	+0.05	-0.03	-0.36
Llama-2-7B-Chat	FP16	8.43	7.85	5.00	2.50	3.10	6.62	8.97	9.85	6.54
	GPTQ	-0.43	-0.55	-0.50	+0.10	-0.65	-1.27	-0.82	-1.42	-0.69
	AWQ	+0.87	-0.45	+0.05	+0.30	-0.60	+0.63	-0.02	-0.15	+0.08
Llama-2-13B-Chat	FP16	8.90	7.80	6.15	2.50	4.20	8.10	8.88	9.85	7.05
	GPTQ	+0.50	-0.50	-0.90	+0.65	-0.75	-1.00	+0.30	0.00	-0.22
	AWQ	-0.12	-0.25	0.00	+0.50	-0.60	-0.75	+0.05	+0.05	-0.14
Llama-2-70B-Chat	FP16	9.25	7.80	6.35	4.05	3.74	7.40	8.95	9.90	7.18
	GPTQ	+0.05	+0.03	+0.10	-0.80	+0.06	0.00	-0.25	0.00	-0.10
	AWQ	-0.18	+0.80	-0.95	+0.15	-0.34	+0.20	-0.40	-0.28	-0.13
Llama-3.1-8B-it	FP16	9.35	9.15	4.60	6.60	6.40	8.62	8.45	9.95	7.89
	GPTQ	-0.25	-1.25	-0.05	+0.50	-1.15	+0.03	+0.73	-0.55	-0.25
	AWQ	-0.60	+0.10	+1.45	-0.95	+0.05	-0.82	+1.17	0.00	+0.05
Llama-3.1-70B-it	FP16	9.65	9.50	7.58	7.65	8.14	9.70	9.93	9.99	9.02
	GPTQ	-0.15	-0.70	+0.22	-0.35	+0.16	+0.05	-0.13	-0.02	-0.12
	AWQ	+0.10	0.00	+0.67	+0.10	-0.19	-0.20	+0.02	-0.04	+0.06
Llama-3.1-405B-it	FP16	9.80	9.55	8.00	7.15	8.60	10.00	9.95	9.93	9.12
	GPTQ	-0.12	0.00	+0.30	-0.23	-0.06	-0.10	0.00	+0.02	-0.02
	AWQ	0.00	-0.05	+0.50	-0.30	+0.20	-0.30	0.00	+0.07	+0.02
Llama-3.2-1B-it	FP16	7.05	6.35	2.40	3.80	3.40	4.15	6.15	8.50	5.22
	GPTQ	-1.30	-0.30	+1.00	-1.10	-0.24	+1.10	-0.50	-1.45	-0.34
	AWQ	-0.35	+0.70	+0.95	-0.32	-0.25	-0.60	-0.07	-0.50	-0.05
Llama-3.2-3B-it	FP16	8.90	8.40	3.90	5.25	5.85	8.10	7.80	8.65	7.11
	GPTQ	-0.70	-0.50	+0.55	-0.05	-0.73	-0.75	-1.00	+0.10	-0.39
	AWQ	-0.05	-0.45	+1.65	+0.40	-0.40	+0.05	-0.25	-0.45	+0.06
Llama-3.3-70B-it	FP16	9.00	9.60	8.45	8.60	8.32	9.65	9.45	10.00	9.13
	GPTQ	0.00	0.00	+0.60	-0.65	-0.62	-0.20	+0.20	-0.10	-0.09
	AWQ	+0.05	-0.30	+0.10	-0.25	-1.02	-0.05	-0.05	-0.05	-0.19

Table 6: For each category, the MT-Bench scores (FP16, GPTQ, AWQ) are presented for 12 models. The difference in MT-Bench scores compared to FP16 is computed and shown for GPTQ and AWQ. The three categories with the largest score drops compared to FP16 are highlighted in red, while The categories with score increases are highlighted in green.

Model	Turn	FP16	GPTQ* (W4/A16)	AWQ (W4/A16)	API Cost
Vicuna-7B-v1.3	1	6.69	6.11	6.56	\$4.96
	2	5.30	5.30	5.50	
	Avg.	6.00	5.71	6.03	
Gemma-2B-it	1	5.09	5.03	4.68	\$4.47
	2	4.18	4.03	4.19	
	Avg.	4.63	4.53	4.43	
Gemma-7B-it	1	6.89	6.40	6.34	\$4.47
	2	5.09	4.73	4.93	
	Avg.	5.99	5.56	5.63	
Llama-2-7B-Chat	1	7.12	6.67	6.86	\$5.80
	2	5.96	4.99	6.38	
	Avg.	6.54	5.85	6.62	
Llama-2-13B-Chat	1	7.34	7.29	7.38	\$5.85
	2	6.75	6.38	6.44	
	Avg.	7.05	6.83	6.91	
Llama-2-70B-Chat	1	7.28	7.21	7.29	\$7.47
	2	7.13	6.95	6.83	
	Avg.	7.18	7.08	7.05	
Llama-3.1-8B-it	1	8.17	7.89	8.52	\$5.49
	2	7.61	7.39	7.36	
	Avg.	7.89	7.64	7.94	
Llama-3.1-70B-it	1	9.24	9.04	9.25	\$5.87
	2	8.80	8.76	8.90	
	Avg.	9.02	8.90	9.08	
Llama-3.1-405B-it	1	9.29	9.28	9.24	\$5.56
	2	8.95	8.88	9.04	
	Avg.	9.12	9.10	9.14	
Llama-3.2-1B-it	1	5.71	5.50	5.56	\$5.06
	2	4.74	4.27	4.78	
	Avg.	5.22	4.88	5.17	
Llama-3.2-3B-it	1	7.58	7.29	7.46	\$4.99
	2	6.64	6.15	6.88	
	Avg.	7.11	6.72	7.17	
Llama-3.3-70B-it	1	9.14	9.13	9.01	\$5.70
	2	9.13	8.95	8.86	
	Avg.	9.13	9.04	8.94	

Table 7: A evaluation of the influence of quantization on the MT-Bench multi-turn conversation benchmark. * denotes the use of AutoGPTQ for quantization. The API cost represents the GPT4 usage fee incurred when running MT-Bench for a single model. We highlighted the best value in green.

Model	Turn	FP16	GPTQ (W4/A16)	AWQ (W4/A16)	Cost
Llama-2-7B	1	1.25	1.29	1.23	9.24 × 3
	2	1.06	1.00	1.06	
	Avg.	1.16	1.14	1.14	
Llama-2-13B	1	1.41	1.45	1.39	9.24 × 3
	2	1.04	1.01	1.05	
	Avg.	1.23	1.23	1.22	

Table 8: MT-Bench (single mode) results for pretrained models without intruction tuning.

Model	Method (Bits)	Dataset	# Samples	Seq. Len.
Vicuna-7B-v1.3	GPTQ*	Wikitex2	128	2,048
	AWQ	mit-han-lab/pile-val-backup	128	512
Gemma-2B-it	GPTQ*	Wikitex2	128	2,048
	AWQ	mit-han-lab/pile-val-backup	128	512
Gemma-7B-it	GPTQ*	Wikitex2	128	2,048
	AWQ	mit-han-lab/pile-val-backup	128	512
Llama-2-7B-Chat	GPTQ*	Wikitex2	128	2,048
	AWQ	mit-han-lab/pile-val-backup	128	512
Llama-2-13B-Chat	GPTQ*	Wikitex2	128	2,048
	AWQ	mit-han-lab/pile-val-backup	128	512
Llama-2-70B-Chat	GPTQ*	Wikitex2	128	2,048
	AWQ	mit-han-lab/pile-val-backup	128	512
Llama-3.1-8B-it	FP8	HuggingFaceH4/ultrachat-200k	512	4,096
	GPTQ*	Wikitex2	128	2,048
	GPTQ**(4/16)	LLM_compression_calibration	512	8,192
	GPTQ**(8/16)	<i>Random data from vocab of tokenizer</i>	256	8,192
	SmoothQuant	<i>Random data from vocab of tokenizer</i>	256	8,192
	AWQ	mit-han-lab/pile-val-backup	128	512
Llama-3.1-70B-it	FP8	HuggingFaceH4/ultrachat-200k	512	4,096
	GPTQ*	Wikitex2	128	2,048
	GPTQ**(4/16)	LLM_compression_calibration	512	8,192
	GPTQ**(8/16)	LLM_compression_calibration	256	8,192
	SmoothQuant	LLM_compression_calibration	256	8,192
	AWQ	mit-han-lab/pile-val-backup	128	512
Llama-3.1-405B-it	FP8	HuggingFaceH4/ultrachat-200k	512	4,096
	GPTQ**(4/16)	LLM_compression_calibration	512	8,192
	AWQ	mit-han-lab/pile-val-backup	128	512
Llama-3.2-1B-it	FP8	LLM_compression_calibration	512	8,192
	GPTQ***	NeelNanda/pile-10k	128	512
	SmoothQuant	LLM_compression_calibration	512	8,192
	AWQ	NeelNanda/pile-10k	128	512
Llama-3.2-3B-it	FP8	LLM_compression_calibration	512	8,192
	GPTQ***	NeelNanda/pile-10k	128	512
	SmoothQuant	LLM_compression_calibration	512	8,192
	AWQ	NeelNanda/pile-10k	128	512
Llama-3.3-70B-it	GPTQ***	NeelNanda/pile-10k	128	512
	AWQ	mit-han-lab/pile-val-backup	128	512

Table 9: Calibration settings for quantization, including dataset used, sample count, fixed sequence length, and default configurations ensure stable performance across models. * denotes the use of AutoGPTQ for quantization, while ** uses vLLM project’s LLM-Compressor for GPTQ quantization. Also, *** denotes the use of Intel’s AutoRound for GPTQ quantization.