

# Low-Bit Quantization Favors Undertrained LLMs: Scaling Laws for Quantized LLMs with 100T Training Tokens

Xu Ouyang<sup>1,2\*</sup> Tao Ge<sup>2†</sup> Thomas Hartvigsen<sup>1</sup> Zhisong Zhang<sup>2</sup> Haitao Mi<sup>2</sup> Dong Yu<sup>2</sup>  
<sup>1</sup>University of Virginia <sup>2</sup>Tencent AI Lab Seattle  
 ftp8nr@virginia.edu getao@global.tencent.com

## Abstract

We reveal that low-bit quantization favors undertrained large language models (LLMs) by observing that models with larger sizes or fewer training tokens experience less quantization-induced degradation (QiD) when applying low-bit quantization, whereas smaller models with extensive training tokens suffer significant QiD. To gain deeper insights into this trend, we study over 1500 quantized LLM checkpoints of various sizes and at different training levels (undertrained or fully trained) in a controlled setting, deriving scaling laws for understanding the relationship between QiD and factors such as the number of training tokens, model size and bit width.

With the derived scaling laws, we propose a novel perspective that we can use QiD to measure an LLM’s training levels and determine the number of training tokens required for fully training LLMs of various sizes. Moreover, we use the scaling laws to predict the quantization performance of different-sized LLMs trained with **100 trillion** tokens. Our projection shows that the low-bit quantization performance of future models, which are expected to be trained with over 100 trillion tokens, may NOT be desirable. This poses a potential challenge for low-bit quantization in the future and highlights the need for awareness of a model’s training level when evaluating low-bit quantization research. To facilitate future research on this problem, we release all the 1500+ quantized checkpoints used in this work at <https://huggingface.co/Xu-Ouyang>.




Figure 1: Scaling laws for predicting Quantization-induced Degradation (QiD, denoted as  $\Delta_q Loss$ ) in 7B, 70B, and 405B models trained on up to 100 trillion ( $10^{14}$ ) tokens. While low-bit quantization yields acceptable QiD for undertrained LLMs (trained with  $\leq 10^{12}$  tokens), it is predicted to become undesirable when applied to fully trained LLMs (e.g., trained with 100 trillion tokens, a milestone expected to be reached in the next few years), particularly for smaller models. Note that the gray areas in this figure indicate levels of QiD that degrade the model’s predictions to a level worse than random guessing.

## 1 Introduction

Quantization (Jacob et al., 2018; Krishnamoorthi, 2018; Banner et al., 2019; Frantar et al., 2022; Shen et al., 2024; Lin et al., 2024; Zhang et al., 2024) is one of the most popular techniques for efficiently deploying large language models (LLMs) by reducing the model’s disk size, memory footprint, and improving inference efficiency through lower precision weights and activations. As model sizes have continued to grow over the past years, researchers have moved beyond conventional 8-bit quantization (Zafir et al., 2019; Dettmers et al., 2022; Zhong et al., 2024) and begun

\*Work done while interning at Tencent AI Lab Seattle.

†Corresponding author




Figure 2: Performance of LLMs after low-bit quantization at different sizes and training levels. It is obvious that the models which are smaller or trained with more tokens suffer from greater quantization-induced degradation.

exploring even lower bit width (Bai et al., 2020; Zhang et al., 2020; Wang et al., 2023; Liu et al., 2023; Egiazarian et al., 2024; Liu et al., 2024; Huang et al., 2024), sparking a surge of research interest in low-bit quantization.

While low-bit quantization works well on some LLM checkpoints with very little quantization-induced degradation (QiD), we have observed that these checkpoints are typically with either larger model sizes or fewer training tokens. In contrast, smaller models or those trained with much more tokens tend to suffer from significant QiD when low-bit quantization is applied. As shown in Figure 2(right), 3-bit quantization results in negligible QiD for a 12 billion parameter LLM up to  $10^{11}$  training tokens, but beyond this point, QiD begins to become pronounced; For smaller models (e.g., 160 million and 1 billion parameters), QiD degradation occurs much earlier and is more severe. With even more extreme 2-bit quantization as shown in Figure 2(left), the trend is similar, but QiD worsens sooner and more significantly. This observation suggests that low-bit quantization tends to favor undertrained LLMs and is less compatible with fully trained LLMs.

To gain deeper insights into this trend, we study over 1500 quantized LLM checkpoints of various sizes (ranging from 160M to 12B) and at different training levels<sup>3</sup> (trained with from 1B to 206B training tokens), analyzing the impact of low-bit quantization on them in a controlled setting. We derive scaling laws to model QiD with respect to the number of training tokens, model size, bit width. According to the derived scaling laws, we propose a novel perspective that we can use QiD to measure an LLM’s training levels and determine the number of training tokens required for fully training an LLM given its size. Moreover, we use the scaling laws to predict the performance of different-sized LLMs with 100 trillion training tokens when applying low-bit quantization. Our projection shows that low-bit quantization of future models, which are expected to be trained with over 100 trillion tokens, may not be desirable, which indicates a potential challenge for low-bit quantization in the future and suggests that a model’s training level should be considered in the evaluation of future low-bit quantization research.

The contributions of this work are threefold:

- We reveal that low-bit quantization favors undertrained LLMs but suffers from significant quantization-induced degradation (QiD) when applied to fully trained LLMs. This insight has been largely overlooked in previous low-bit quantization research: very few studies have considered the training level of a quantized LLM when evaluating their proposed low-bit quantization approaches.
- We derive scaling laws to model QiD with respect to the number of training tokens, model size and bit width. Using these scaling laws, we propose to use QiD as a signal to measure whether an LLM is fully trained and estimate the number of training tokens required for LLMs of different sizes to reach a fully trained state. Moreover, we use the scaling law to predict the performance of low-bit quantization for different-sized LLMs trained with 100 trillion tokens. Our projection indicates potential challenges for the future application of low-bit quantization.
- We release all the 1500+ quantized checkpoints used in this work to facilitate future research on this problem.

## 2 Preliminary: Scaling Laws for Large Language Models

Scaling laws for large language models (Kaplan et al., 2020; Hoffmann et al., 2022) are crucial for understanding how these models’ performance improves with increased scale, including the number of parameters and training tokens:

<sup>3</sup>Training levels in this work refer to the extent to which an LLM has been trained (e.g., undertrained, fully trained, or overtrained), which are related to both the number of training tokens and the model size.

**Number of Parameters** LLMs’ performance typically follows a power-law improvement as the number of parameters increases, allowing larger models to better fit and generalize on the same dataset:

$$L(N) = \frac{a}{N^\alpha} + \epsilon \quad (1)$$

where  $L(N)$  is the loss function<sup>4</sup> dependent on  $N$  (the number of non-embedding parameters),  $a$  is a constant (i.e., coefficient),  $\alpha$  is the scaling exponent, and  $\epsilon$  represents the error term. This relationship indicates larger models are generally more capable of capturing the complexities of language, leading to better generalization and lower loss.

**Training Tokens** More training tokens also boost performance in a power-law fashion, enabling models to capture language complexities more effectively:

$$L(D) = \frac{b}{D^\beta} + \epsilon \quad (2)$$

where  $D$  denotes the number of training tokens,  $b$  is a constant (i.e., coefficient) and  $\beta$  is the scaling exponent for training tokens. More training tokens enhance an LLM’s ability to learn and generalize, allowing it to achieve better language modeling performance with lower loss.

When scaling both the number of parameters  $N$  and the amount of training data  $D$  simultaneously, the scaling law can be expressed as a function that accounts for the combined effects of both:

$$L(N, D) = \left[ \left( \frac{N_c}{N} \right)^{\frac{\alpha_N}{\alpha_D}} + \frac{D_c}{D} \right]^{\alpha_D} \quad (3)$$

This scaling law allows us to estimate the performance of language models at unprecedented scales of model size and training data effectively before conducting actual training runs.

### 3 Scaling Laws for Low-bit Quantization

In this section, we propose scaling laws for low-bit quantization. Unlike the scaling laws discussed in Section 2, the focus here is on understanding how quantization-induced degradation (QiD) changes when low-bit quantization is applied to LLMs of varying training scales. Formally, QiD is defined as follows:

$$\Delta_q Loss = Loss_q - Loss_{16\text{-bit}} \quad (4)$$

where  $Loss_q$  is the cross-entropy loss of a quantized LLM, and  $Loss_{16\text{-bit}}$  is the cross-entropy loss of its pre-quantized counterpart with fp16 or bf16 weights.  $\Delta_q Loss$  represents QiD, which is the difference in loss before and after applying low-bit quantization.

Inspired by conventional scaling laws for language modeling, we investigate the impact of model size and the number of training tokens on QiD. Additionally, we consider bit width (i.e., the precision of quantized weight values).

#### 3.1 Experimental Setting

We select open-sourced LLMs from the Pythia suite (Biderman et al., 2023) for our experiments. Pythia not only includes LLMs of various sizes, but also provides access to all checkpoints throughout its training process (from scratch to 300 billion tokens), allowing us to conduct experiments in a controlled setting to derive scaling laws for low-bit quantization.

We choose 6 different sizes of Pythia LLMs: 160M, 410M, 1B, 2.8B, 6.9B, and 12B. For each size, we sample 20 checkpoints (see Appendix A.1) up to 98k steps<sup>5</sup>.

For quantization, we employ one of the most popular LLM quantization techniques – GPTQ (Frantar et al., 2022) – to quantize the Pythia checkpoints to 2-bit, 3-bit and 4-bit levels.

We evaluate QiD on 1,000 randomly sampled texts from RefinedWeb dataset (Penedo et al., 2023).

#### 3.2 Training Tokens

In contrast to traditional language modeling scaling laws where the number of training tokens  $D$  appears in the denominator, we propose the relationship between training tokens and QiD as follows:

$$\Delta_q Loss(D) \approx b \cdot D^\beta \quad (5)$$

because the more training tokens, the more significant the QiD becomes, according to our observations in Figure 2.

<sup>4</sup>We mainly discuss cross entropy loss for language modeling in this paper.

<sup>5</sup>98k steps correspond to approximately 206 billion tokens, which is equivalent to one epoch of Pythia’s training data. Although Pythia was trained for 143k steps, we skipped checkpoints beyond 98k steps to avoid the influence of duplicated data, as the data beyond 98k steps probably represents the second epoch with data that has already been trained with.




Figure 3: The fitted scaling law of QiD with respect to the number of training tokens in the form of Eq (5), where  $\beta$  is fitted to be 0.5316.

We use the above functional form to fit the QiD observed in the quantized Pythia checkpoints in 3, obtaining  $\beta = 0.5316$ , which fits the trend of QiD with respect to the change in training tokens quite well.

### 3.3 Model Size

As mentioned in Figure 2, the larger the size of the model, the smaller the QiD tends to be. Therefore, we propose the relationship between model size (i.e., the number of non-embedding parameters) and QiD as follows:

$$\Delta_q Loss(N) \approx \frac{a}{N^\alpha} \quad (6)$$

We use the above functional form to fit the QiD of quantized Pythia checkpoints in Figure 4, obtaining  $\alpha = 0.2276$ .




Figure 4: The fitted scaling law of QiD with respect to the model size (i.e., the number of non-embedding parameters) in the form of Eq (6), where  $\alpha$  is fitted to be 0.2276.




Figure 5: The fitted scaling law of QiD with respect to the bit width in the form of Eq (7), where  $\gamma$  is fitted to be 5.4812.

### 3.4 Bit Width

Bit width is a factor not present in conventional scaling laws. Considering that the role of bit width is similar to that of the number of parameters (both aim to increase the model’s expressiveness), we propose a similar functional form as in Section 3.3 to model bit width in Eq (7), and fit the data points of Pythia in Figure 5:

$$\Delta_q Loss(P) \approx \frac{c}{P^\gamma} \quad (7)$$

### 3.5 Unified Scaling Law

With the basic scaling laws derived in Sections 3.2 (the number of training tokens), 3.3 (model size), and 3.4 (bit width), we study how to model QiD with all three factors together. Inspired by Kaplan et al. (2020), we consider the following four principles for unifying the factors:

- Fixing  $D$  and  $P$ , sending  $N \rightarrow \infty$ , we expect  $\Delta_q Loss \rightarrow 0$ .
- Fixing  $N$  and  $P$ , sending  $D \rightarrow 0$ , we expect  $\Delta_q Loss \rightarrow 0$ .
- Fixing  $N$  and  $D$ , sending  $P \geq 16$ , we expect  $\Delta_q Loss \rightarrow 0$ .
- Fixing  $N$  and  $D$ , sending  $P \rightarrow 0$ ,  $\Delta_q Loss$  should be very large.

We propose the unified scaling law for low-bit quantization as follows:

$$\Delta_q Loss(N, D, P) = k \cdot \frac{D^\beta}{N^\alpha P^\gamma} \quad (8)$$

where  $k$  is the joint coefficient, and both the coefficient and exponents ( $\alpha, \beta, \gamma$ ) are positive. Figure 6 displays the fitted curves using this functional form. The jointly fitted exponents  $\alpha, \beta$ , and  $\gamma$  closely match those obtained by fitting these variables independently, further validating the effectiveness of the joint function form  $\Delta_q Loss(N, D, P)$ .




Figure 6: The unified scaling law we fit based on Eq (8) with the GPTQ-quantized LLMs from the Pythia suite:  $\Delta_q Loss(N, D, P) = 0.017D^{0.5251}/(N^{0.2261} \cdot P^{5.4967})$

Given the unified scaling law for  $\Delta_q Loss$  and the definition of  $\Delta_q Loss$  in Eq (4), we can easily predict a quantized LLM’s performance as  $Loss_q = Loss_{16-bit} + \Delta_q Loss$ , as illustrated in Figure 7, which fits well with the observed data points.




Figure 7: We can predict the performance of a quantized LLM as  $Loss_q = Loss_{16-bit} + \Delta_q Loss$ , where  $Loss_{16-bit}$  can be predicted by the conventional LLM’s scaling law which is fitted based on the function form of Eq (3) with the LLMs in the Pythia suite as  $Loss_{16-bit} = [(4.74e^{19}/N)^{(0.045/0.399)} + 7.63e^{10}/D]^{0.399}$ .

### 3.6 Validation with Ablation Studies

We validate the scaling law derived in Section 3.5 with different test data, quantization methods and foundation models.

#### 3.6.1 Test Data

We compare the results obtained using RefinedWeb and Wikitext-2 (Merity et al., 2016) as test data in Figure 8, demonstrating that the QiD results on these two test datasets are almost identical. This suggests that the trends of QiD are largely independent of the test data.




Figure 8: QiD results evaluated on RefinedWeb and Wikitext-2 with the 12B Pythia model.

### 3.6.2 Quantization Methods

We quantize the Pythia checkpoints using another two popular quantization methods – AWQ (Lin et al., 2024) and bitsandbytes<sup>6</sup> in addition to GPTQ. We show the QiD results and fitted scaling laws in Figure 9, and we observe that the QiD trends for different quantization methods are almost identical, although the fitted scaling laws show slight differences.




Figure 9: QiD results and fitted scaling laws for different quantization methods. Note that the GPTQ function here differs slightly from that in Figure 6, as it is fitted exclusively with 4-bit quantized Pythia checkpoints, whereas the function in Figure 6 is fitted using all quantized Pythia checkpoints.

### 3.6.3 Foundation Models




Figure 10: **Left:** Scaling laws for low-bit quantization, fitted on the LLM checkpoints of the Spectra suite, which are all trained with 300B tokens; **Right:** Actual  $\Delta_q Loss$  VS Predicted  $\Delta_q Loss$  that is computed based on the scaling laws fitted on Llama and Qwen.

Figure 10 shows the fitting results of our scaling laws function form, Eq (8), on the Spectra suite (Kaushal et al., 2024) as well as the popular open-sourced Llama (Touvron et al., 2023; Dubey et al., 2024) and Qwen (Yang et al., 2024) models, which confirms that the scaling laws are not only valid for Pythia but are likely to be broadly applicable.

<sup>6</sup><https://github.com/bitsandbytes-foundation/bitsandbytes>

## 4 Discussion: Low-bit Quantization Favors Undertrained LLMs

### 4.1 Intuition

Based on the scaling laws we derived in Section 3, we confirm low-bit quantization tends to favor models with fewer training tokens or larger model sizes, which are essentially undertrained LLMs.




Figure 11: Fully trained LLMs suffer from much greater QiD (i.e.,  $\Delta_q Loss$ ) than undertrained LLMs.

Figure 11 illustrates the relationship between QiD, model size, and training tokens. Points located in the upper-left corner are more fully trained and have a much higher QiD, while points in the bottom-right corner are more undertrained and have a lower QiD.




Figure 12: Changes in model weights between adjacent checkpoints. Early (undertrained) checkpoints exhibit significant weight fluctuations during training, making the model relatively robust to weight variations. Therefore, small changes introduced by quantization have a limited impact on the model’s performance. In contrast, fully trained checkpoints demonstrate very little weight fluctuations during training. As a result, low-bit quantization is likely to push weights **beyond the narrow range of recent variations**, leading to performance degradation or even model collapse.

To understand this observation intuitively, we illustrate changes in sampled model weights between adjacent checkpoints in Figure 12. It can be observed that the early checkpoints exhibit significant changes in weights. Due to the significant fluctuations in weights during training, the model becomes inherently robust to weight variations, meaning that even if low-bit quantization introduces some precision loss, the overall impact on the model remains limited. On the other hand, checkpoints from the later stages of training, which are more fully trained, show very small changes in weights (often at a very small scale, even beyond the 3rd-4th decimal place). In such cases, low-bit quantization is very likely to shift weights outside the small range of recent variations, potentially causing the model to degrade or even collapse.

From another perspective, during the undertrained stage, the model’s weights undergo significant changes and have not yet fully exploited the precision dimension. In the later, more fully trained stage, as weight adjustments stabilize, the model increasingly relies on precision to continue optimizing the training objective and improving language modeling performance. This aligns with the two phases of representation learning in the information bottleneck theory (Shwartz-Ziv & Tishby, 2017): during the early training phase, gradients have a large mean and small variance, making high precision unnecessary. However, in the later training phase, gradients have a small mean and large variance, requiring higher precision for the model to converge effectively.

Table 1: Prediction of the number of training tokens (in trillion) needed to achieve a given training level measured by  $\Delta_q Loss$  for different model sizes and bit widths. Note that  $\Delta_q Loss = 0.2$  means the likelihood is reduced to 80% of its original value ( $e^{-0.2} \approx 0.8$ ), while  $\Delta_q Loss = 0.5$  means the likelihood is reduced to 60% ( $e^{-0.5} \approx 0.6$ ).

Model Size	$\Delta_q Loss = 0.2$			$\Delta_q Loss = 0.3$			$\Delta_q Loss = 0.4$			$\Delta_q Loss = 0.5$		
	2 bits	3 bits	4 bits	2 bits	3 bits	4 bits	2 bits	3 bits	4 bits	2 bits	3 bits	4 bits
1B	0.0011	0.1089	1.4424	0.0025	0.1990	2.6786	0.0043	0.3051	4.1556	0.0066	0.4251	5.8422
7B	0.0026	0.3038	4.5066	0.0057	0.5550	8.3689	0.0099	0.8512	12.9836	0.0152	1.1860	18.2531
70B	0.0071	1.0228	<b>17.3499</b>	0.0154	1.8687	32.2192	0.0267	2.8659	49.9854	0.0409	3.9932	70.2723
405B	0.0151	2.5807	<b>48.4861</b>	0.0328	4.7151	90.0398	0.0567	7.2311	139.6892	0.0868	10.0754	196.3829

## 4.2 QiD: A Signal that Measures an LLM’s Training Level

Unlike previous work that often uses the inability of the loss to decrease further as a signal to determine whether an LLM is fully trained (i.e., saturated), we introduce a novel perspective that we can use QiD to determine whether an LLM is fully trained. If an LLM exhibits  $QiD \approx 0$  after low-bit quantization, it suggests that the LLM is likely undertrained, as it has not yet exploited higher precision, as discussed in Section 4.1.

With the scaling law in Eq (8) derived in Section 3.5, we can estimate how many training tokens are needed for a given LLM size to be considered fully trained based on QiD predictions. Table 1 shows the number of training tokens required for different model sizes to achieve  $\Delta_q Loss = \{0.2, 0.3, 0.4, 0.5\}$  when applying low-bit quantization. For a 70B scale model, achieving a QiD greater than 0.2 (corresponding to likelihood decrease by 20%) under 4-bit quantization requires over 17 trillion training tokens. In contrast, for a 405B scale LLM, achieving a QiD above 0.2 under 4-bit quantization requires nearly 50 trillion training tokens – a scale far beyond what has been achieved by now, indicating that current training efforts for extremely large LLMs may be still far from sufficient.




Figure 13: The number of training tokens for the state-of-the-art 7B-scale LLMs increase by nearly 50 $\times$  over the past 4 years. According to this trend, it is expected that the future models will have much more training tokens.

## 4.3 QiD Prediction When Scaling to 100 Trillion Training Tokens

Figure 13 shows the trend in the number of training tokens for state-of-the-art 7B-scale LLMs from 2020 to the present, showing that the number of training tokens has increased nearly 50 $\times$  over the past 4 years. Based on this trend, it is very likely that LLMs in 2025-2026 will be trained with up to 100 trillion ( $10^{14}$ ) tokens<sup>7</sup>.

Using the scaling laws derived, we predict the performance of quantized LLMs trained on 100 trillion tokens, as illustrated in Figure 1 at the beginning of this paper. In particular, performance degradation with 2-bit and 3-bit quantization at the unprecedented training scale of 100 trillion tokens is predicted to be severe, which is in stark contrast to the acceptable performance at the current training scale of  $10^{13}$  tokens. This indicates a challenge for the practical application of low-bit quantization to future LLMs.

<sup>7</sup>Although there have been claims that internet data is nearing exhaustion, recent continuous innovations in synthetic data creation (Ge et al., 2024) lead us to believe that the milestone of 100 trillion training tokens is achievable in the next few years.

#### 4.4 From Low-bit Quantization to Low-bit LLMs

Although this work mainly focuses on low-bit (post-)quantization, we suspect that native low-bit LLMs are also likely to favor undertrained LLMs. We replicated the popular 1-bit LLM – *BitNet b1.58* (Ma et al., 2024) – to compare it with its bf16 counterpart throughout training. Specifically, we trained 120M and 1.2B decoder-only models with both bf16 and BitNet. Figure 14 shows the comparison of training losses between BitNet and its 16-bit counterparts in the early- and mid-training steps. It can be observed that, in the early stages of training, the training loss curves of BitNet closely match (and even outperform) those of bf16, as BitNet tends to use a higher learning rate than bf16 training according to its training recipe. As training continues, the 120M BitNet gradually begins to lag behind its bf16 counterpart, and after further training steps, a noticeable gap starts to appear in the 1.2B models, which is consistent with our observations for low-bit quantization. This indicates that native low-bit LLMs such as BitNet<sup>8</sup> may also favor undertrained LLMs, though the gap manifests later compared to post-quantization, as native low-bit training keeps the model capable of operating under low precision throughout the training process.




Figure 14: Training losses of BitNet and its 16-bit counterparts show a trend similar to that of low-bit quantization – they tend to perform well when undertrained but struggle to match the performance of fully trained LLMs.

## 5 Conclusion

We derive scaling laws for low-bit quantization from over 1500 quantized LLM checkpoints, and reveal that low-bit quantization favors undertrained LLMs. We provide an intuitive interpretation for this phenomenon and introduce a novel perspective of using QiD as a signal to determine a model’s training level. Moreover, we use the derived scaling laws to predict the effect of low-bit quantization on LLMs trained with 100 trillion tokens. This, on one hand, challenges the future practical value of low-bit quantization, and on the other hand, suggests that future research on low-bit quantization should consider the model’s training level during evaluation. Alongside concurrent research (Kumar et al., 2024; Feng et al., 2024) that takes a serious look at the limits of low-bit LLMs, we hope this work can help the community cool down from the surrounding hype, and foster deeper reflection and critical examination in this field.

## Limitations

This work includes the following limitations:

- Although we have done our best to conduct extensive experiments and derive the scaling laws from over 1500 quantized checkpoints, it is still not extensive enough. For example, the training tokens used in our experiments with Pythia only amount to 300 billion. We expect more observations from a greater number of quantized checkpoints in the future to refine the scaling laws we have derived.
- The scaling laws derived in this work are primarily focused on single-stage pre-trained language models. However, advanced LLMs today often employ multi-stage training strategies including supervised fine-tuning and preference optimization, and even within pre-training, multiple stages are often involved (e.g., Llama-3.1 focuses more on high-quality text, math, reasoning, and code data during the final pre-training stages). Such multi-stage training strategies may cause the behavior of the model after quantization to be significantly different, which we plan to explore in future work.

<sup>8</sup>We reviewed the original BitNet paper and some open-sourced reimplementations, and found that their numbers of training tokens were up to 100 billion. Considering their model sizes and the fact that the performance gap of native low-bit LLMs tends to emerge later compared to post-quantization, we express concerns about their performance at larger training scales (i.e., with more training tokens). We call for results of native low-bit LLMs at larger training scales to better justify their practical value.

## References

- Haoli Bai, Wei Zhang, Lu Hou, Lifeng Shang, Jing Jin, Xin Jiang, Qun Liu, Michael Lyu, and Irwin King. Binarybert: Pushing the limit of bert quantization. *arXiv preprint arXiv:2012.15701*, 2020.
- Ron Banner, Yury Nahshan, and Daniel Soudry. Post training 4-bit quantization of convolutional networks for rapid-deployment. *Advances in Neural Information Processing Systems*, 32, 2019.
- Stella Biderman, Hailey Schoelkopf, Quentin Gregory Anthony, Herbie Bradley, Kyle O’Brien, Eric Hallahan, Mohammad Aflah Khan, Shivanshu Purohit, USVSN Sai Prashanth, Edward Raff, et al. Pythia: A suite for analyzing large language models across training and scaling. In *International Conference on Machine Learning*, pp. 2397–2430. PMLR, 2023.
- 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.
- 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.
- Vage Egiazarian, Andrei Panferov, Denis Kuznedelev, Elias Frantar, Artem Babenko, and Dan Alistarh. Extreme compression of large language models via additive quantization. *arXiv preprint arXiv:2401.06118*, 2024.
- Guhao Feng, Kai Yang, Yuntian Gu, Xinyue Ai, Shengjie Luo, Jiacheng Sun, Di He, Zhenguo Li, and Liwei Wang. How numerical precision affects mathematical reasoning capabilities of llms. *arXiv preprint arXiv:2410.13857*, 2024.
- Elias Frantar, Saleh Ashkboos, Torsten Hoefler, and Dan Alistarh. Gptq: Accurate post-training quantization for generative pre-trained transformers. *arXiv preprint arXiv:2210.17323*, 2022.
- Tao Ge, Xin Chan, Xiaoyang Wang, Dian Yu, Haitao Mi, and Dong Yu. Scaling synthetic data creation with 1,000,000,000 personas. *arXiv preprint arXiv:2406.20094*, 2024.
- Jordan Hoffmann, Sebastian Borgeaud, Arthur Mensch, Elena Buchatskaya, Trevor Cai, Eliza Rutherford, Diego de Las Casas, Lisa Anne Hendricks, Johannes Welbl, Aidan Clark, et al. Training compute-optimal large language models. *arXiv preprint arXiv:2203.15556*, 2022.
- Wei Huang, Yangdong Liu, Haotong Qin, Ying Li, Shiming Zhang, Xianglong Liu, Michele Magno, and Xiaojuan Qi. Billm: Pushing the limit of post-training quantization for llms. *arXiv preprint arXiv:2402.04291*, 2024.
- Benoit Jacob, Skirmantas Kligys, Bo Chen, Menglong Zhu, Matthew Tang, Andrew Howard, Hartwig Adam, and Dmitry Kalenichenko. Quantization and training of neural networks for efficient integer-arithmetic-only inference. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 2704–2713, 2018.
- Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. Scaling laws for neural language models. *arXiv preprint arXiv:2001.08361*, 2020.
- Ayush Kaushal, Tejas Vaidhya, Arnab Kumar Mondal, Tejas Pandey, Aaryan Bhagat, and Irina Rish. Spectra: Surprising effectiveness of pretraining ternary language models at scale. *arXiv preprint arXiv:2407.12327*, 2024.
- Raghuraman Krishnamoorthi. Quantizing deep convolutional networks for efficient inference: A whitepaper. *arXiv preprint arXiv:1806.08342*, 2018.
- Tanishq Kumar, Zachary Ankner, Benjamin F Spector, Blake Bordelon, Niklas Muennighoff, Mansheej Paul, Cengiz Pehlevan, Christopher Ré, and Aditi Raghunathan. Scaling laws for precision. *arXiv preprint arXiv:2411.04330*, 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.
- Jing Liu, Ruihao Gong, Xiuying Wei, Zhiwei Dong, Jianfei Cai, and Bohan Zhuang. Qllm: Accurate and efficient low-bitwidth quantization for large language models. *arXiv preprint arXiv:2310.08041*, 2023.
- Zirui Liu, Jiayi Yuan, Hongye Jin, Shaochen Zhong, Zhaozhuo Xu, Vladimir Braverman, Beidi Chen, and Xia Hu. Kivi: A tuning-free asymmetric 2bit quantization for kv cache. *arXiv preprint arXiv:2402.02750*, 2024.

- Shuming Ma, Hongyu Wang, Lingxiao Ma, Lei Wang, Wenhui Wang, Shaohan Huang, Li Dong, Ruiping Wang, Jilong Xue, and Furu Wei. The era of 1-bit llms: All large language models are in 1.58 bits. *arXiv preprint arXiv:2402.17764*, 2024.
- Stephen Merity, Caiming Xiong, James Bradbury, and Richard Socher. Pointer sentinel mixture models. *arXiv preprint arXiv:1609.07843*, 2016.
- Guilherme Penedo, Quentin Malartic, Daniel Hesslow, Ruxandra Cojocaru, Alessandro Cappelli, Hamza Alobeidli, Baptiste Pannier, Ebtesam Almazrouei, and Julien Launay. The refinedweb dataset for falcon llm: outperforming curated corpora with web data, and web data only. *arXiv preprint arXiv:2306.01116*, 2023.
- Xuan Shen, Peiyan Dong, Lei Lu, Zhenglun Kong, Zhengang Li, Ming Lin, Chao Wu, and Yanzhi Wang. Agile-quant: Activation-guided quantization for faster inference of llms on the edge. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pp. 18944–18951, 2024.
- Ravid Shwartz-Ziv and Naftali Tishby. Opening the black box of deep neural networks via information. *arXiv preprint arXiv:1703.00810*, 2017.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shriti Bhosale, et al. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*, 2023.
- Hongyu Wang, Shuming Ma, Li Dong, Shaohan Huang, Huaijie Wang, Lingxiao Ma, Fan Yang, Ruiping Wang, Yi Wu, and Furu Wei. Bitnet: Scaling 1-bit transformers for large language models. *arXiv preprint arXiv:2310.11453*, 2023.
- An Yang, Baosong Yang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Zhou, Chengpeng Li, Chengyuan Li, Dayiheng Liu, Fei Huang, et al. Qwen2 technical report. *arXiv preprint arXiv:2407.10671*, 2024.
- Ofir Zafrir, Guy Boudoukh, Peter Izsak, and Moshe Wasserblat. Q8bert: Quantized 8bit bert. In *2019 Fifth Workshop on Energy Efficient Machine Learning and Cognitive Computing-NeurIPS Edition (EMC2-NIPS)*, pp. 36–39. IEEE, 2019.
- Cheng Zhang, Jianyi Cheng, George A Constantinides, and Yiren Zhao. Lqer: Low-rank quantization error reconstruction for llms. *arXiv preprint arXiv:2402.02446*, 2024.
- Wei Zhang, Lu Hou, Yichun Yin, Lifeng Shang, Xiao Chen, Xin Jiang, and Qun Liu. Ternarybert: Distillation-aware ultra-low bit bert. *arXiv preprint arXiv:2009.12812*, 2020.
- Yunshan Zhong, Jiawei Hu, You Huang, Yuxin Zhang, and Rongrong Ji. Erq: Error reduction for post-training quantization of vision transformers. *arXiv preprint arXiv:2407.06794*, 2024.

## A Appendix

### A.1 Implementation Details

**Checkpoints of the Pythia** We choose the following 20 checkpoints of the Pythia models at the following steps for fitting the scaling laws: {512, 1k, 2k, 4k, 6k, 8k, 10k, 12k, 14k, 20k, 24k, 29k, 36k, 43k, 57k, 71k, 86k, 93k, 95k, 98k}.

**Tokenization consistency** To ensure consistency in token counts for computing cross entropy loss, which can vary with different tokenizers, we use the token counts generated by the Llama-3 8B (Dubey et al., 2024) tokenizer for all QiD calculations in this work.