
CLoQ: Enhancing Fine-Tuning of Quantized LLMs via Calibrated LoRA Initialization

Yanxia Deng^{*1} Aozhong Zhang^{*1} Naigang Wang² Selcuk Gurses¹ Zi Yang¹ Penghang Yin¹

Abstract

Fine-tuning large language models (LLMs) using low-rank adaptation (LoRA) has become a highly efficient approach for downstream tasks, particularly in scenarios with limited computational resources. However, applying LoRA techniques to quantized LLMs poses unique challenges due to the reduced representational precision of quantized weights. In this paper, we introduce CLoQ (Calibrated LoRA initialization for Quantized LLMs), a simplistic initialization strategy designed to overcome these challenges. Our approach focuses on minimizing the layer-wise discrepancy between the original LLM and its quantized counterpart with LoRA components during initialization. By leveraging a small calibration dataset, CLoQ quantizes a pre-trained LLM and determines the optimal LoRA components for each layer, ensuring a strong foundation for subsequent fine-tuning. A key contribution of this work is a novel theoretical result that enables the accurate and closed-form construction of these optimal LoRA components. We validate the efficacy of CLoQ across multiple tasks such as language generation, arithmetic reasoning, and commonsense reasoning, demonstrating that it consistently outperforms existing LoRA fine-tuning methods for quantized LLMs, especially at ultra low-bit widths. The code is available at <https://github.com/AozhongZhang/CLoQ>

1. Introduction

Large language models (LLMs) (Achiam et al., 2023; Touvron et al., 2023; Jiang et al., 2023; Guo et al., 2024a) have achieved remarkable success across a wide range of domains and applications. With ongoing advancements, the size and complexity of LLMs have grown exponentially, with some models now exceeding billions or even trillions of parameters. Although this scaling has unlocked

unprecedented capabilities, it also introduces significant challenges, particularly in efficiently adapting these models to downstream tasks. Traditionally, full fine-tuning has been the dominant approach for adapting pre-trained models, involving updates to all model parameters. While effective in achieving state-of-the-art results, full fine-tuning is resource-intensive, requiring substantial GPU memory to store both model weights and optimizer states. These memory demands grow with the size of the model, making full fine-tuning increasingly impractical for large-scale models in resource-constrained settings.

To address these challenges, parameter-efficient fine-tuning (PEFT) (Houlsby et al., 2019), such as Low-Rank Adaptation (LoRA) (Hu et al., 2021), has emerged as a promising approach. PEFT updates only a small subset of parameters while keeping the majority of the model unchanged, enabling resource-efficient fine-tuning of large-scale models. LoRA, for instance, introduces small, learnable low-rank matrices into the model’s architecture. These matrices are optimized during fine-tuning while the original model weights remain frozen, significantly reducing memory and computational requirements. This design leverages the insight that weight updates often reside in a low dimensional subspace, allowing LoRA to achieve efficient adaptation with minimal overhead.

In an orthogonal direction, model compression techniques, notably quantization (Rastegari et al., 2016; Hubara et al., 2018; Choi et al., 2018; Wang et al., 2018; Shao et al., 2023; Frantar et al., 2022a; Zhang et al., 2024a), have been developed to minimize GPU memory usage by converting high-precision weights into low-precision representations. Initially designed for inference in memory-limited environments, quantization methods have since been adapted to support fine-tuning. A notable advancement in this regard is QLoRA (Dettmers et al., 2023), which combines LoRA with quantization techniques to reduce GPU memory requirements for fine-tuning significantly. By leveraging low-rank adaptation and quantized weights, QLoRA enables resource-efficient fine-tuning of LLMs without compromising performance, making it a powerful tool for adapting large-scale models to downstream tasks. However, extending LoRA to quantized LLMs introduces additional challenges, as the reduced representational precision of quantized weights can disrupt standard initialization strategies,

^{*}Equal contribution ¹Department of Mathematics and Statistics, University at Albany, SUNY ²IBM T. J. Watson Research Center. Correspondence to: Penghang Yin <pyin@albany.edu>.

impacting task performance. Recent works (Li et al., 2023; Yao et al., 2023; Liao et al., 2024; Guo et al., 2024b) proposed minimizing quantization error through strategic initialization of the low-rank components in LoRA, to align with the original weight states of the model. This strategy has demonstrated success in fine-tuning lower-bit quantized LLMs.

Contributions. In this paper, we introduce CLOQ (Calibrated LoRA for Quantized LLMs), an efficient layer-wise, data-driven initialization strategy specifically designed for quantized LLMs by leveraging a small calibration dataset. CLOQ consists of two main steps: a post-training quantization phase to obtain quantized weights and a generalized low-rank approximation phase under a linear transformation to compute the corresponding optimal adapters. We derive a novel closed-form solution to the low-rank approximation problem, which can be efficiently computed using just two singular value decompositions (SVDs).

Our CLOQ method requires no back-propagation, making it a highly efficient approach for fine-tuning quantized models. We demonstrate the effectiveness of CLOQ through extensive validation on multiple benchmark datasets. Our results show that CLOQ consistently outperforms existing LoRA methods for quantized LLMs, particularly at ultra-low bit-widths. For instance, the fine-tuning accuracy of INT2 CLOQ on the Llama2-13B model (Touvron et al., 2023) surpasses that of INT4 QLoRA (Dettmers et al., 2023) in the arithmetic reasoning tasks, as shown in Table 1.

Notations. We clarify the mathematical notations that will be used throughout this paper: we denote vectors by bold small letters and matrices by bold capital ones. For any matrix $\mathbf{X} \in \mathbb{R}^{m \times n}$, $\mathbf{X}^\top \in \mathbb{R}^{n \times m}$ is the transpose of \mathbf{X} , and $\text{Tr}(\mathbf{X}) := \sum_{i=1}^m X_{i,i}$ denotes the trace of \mathbf{X} when $m = n$. For any two matrices $\mathbf{X}, \mathbf{Y} \in \mathbb{R}^{m \times n}$, $\langle \mathbf{X}, \mathbf{Y} \rangle := \text{Tr}(\mathbf{X}^\top \mathbf{Y}) = \sum_{i=1}^m \sum_{j=1}^n X_{i,j} Y_{i,j}$ is the inner product. We denote the Frobenius norm of \mathbf{X} by $\|\mathbf{X}\|_F := \sqrt{\text{Tr}(\mathbf{X}^\top \mathbf{X})} = \sqrt{\sum_{i=1}^m \sum_{j=1}^n X_{i,j}^2}$. In addition, for any diagonal matrix $\mathbf{\Sigma} = \text{diag}(\sigma_1, \dots, \sigma_n)$, we denote its square root by $\mathbf{\Sigma}^{\frac{1}{2}} := \text{diag}(\sigma_1^{\frac{1}{2}}, \dots, \sigma_n^{\frac{1}{2}})$.

2. Background

Integer Quantizer. Given a set of m weights $\mathbf{w} \in \mathbb{R}^m$, the widely-used b -bit uniform integer (INT) quantizer (Choi et al., 2018) determines the float scaling factor $\delta = \frac{\max(\mathbf{w}) - \min(\mathbf{w})}{2^b - 1}$ and zero-point $z = -\left\lfloor \frac{\min(\mathbf{w})}{\delta} \right\rfloor$, where $\lfloor \cdot \rfloor$ is nearest-to-round operation. The quantizer projects \mathbf{w} onto the equally spaced grids $\mathbb{Q} = \{z \cdot \delta, (z+1) \cdot \delta, \dots, (z + (2^b - 1)) \cdot \delta\}^m$ to obtain the quantized weights

$$\mathbf{q} = \delta \cdot \left(\text{clip} \left(\left\lfloor \frac{\mathbf{w}}{\delta} \right\rfloor + z, 0, 2^b - 1 \right) - z \right).$$

For channel-wise (or group-wise) quantization, the scaling factor δ is shared by the quantized weights within the same channel (or group, respectively).

Post-Training Quantization. Post-training quantization (PTQ) has been a cornerstone technique for compressing LLMs. PTQ methods directly identify low-precision representations of a model without requiring retraining, making them particularly well-suited for extremely large-scale AI models including LLMs. The simplest approach in this category is data-free quantization, commonly referred to as RTN, which involves rounding the pre-trained model weights to their nearest quantized states. More advanced PTQ algorithms, such as OPTQ (Frantar et al., 2022b), leverage a small calibration dataset to solve a least-squares problem under discrete constraints:

$$\min_{\mathbf{Q} \in \mathbb{Q}} \|\mathbf{X}\mathbf{Q} - \mathbf{X}\mathbf{W}\|_F^2, \quad (1)$$

to calibrate the quantization layer by layer. In this formulation, \mathbf{X} denotes the activation or feature matrix corresponding to a batch of calibration data for a fixed layer. This approach ensures that the quantization process preserves the layer’s output by minimizing the discrepancy between the original and quantized representations on the calibration dataset. Several efficient back-propagation-free algorithms (Behdin et al., 2023; Zhang et al., 2024b) have been proposed to address (1).

Low-Rank Adaptation. LoRA (Hu et al., 2021) enables efficient fine-tuning of large pre-trained models by introducing two small, trainable matrices, \mathbf{A} and \mathbf{B} , which are added to the frozen weight matrix \mathbf{W} . The weights of the fine-tuned model are then expressed as $\mathbf{W} + \mathbf{A}\mathbf{B}^\top$, where $\mathbf{A} \in \mathbb{R}^{m \times r}$, $\mathbf{B} \in \mathbb{R}^{n \times r}$, and $r \ll \min(m, n)$. During fine-tuning, only \mathbf{A} and \mathbf{B} are updated, while \mathbf{W} remains fixed, significantly reducing the number of trainable parameters and computational overhead. LoRA initializes its parameters as follows:

$$\mathbf{A} \sim \mathcal{N}(0, \sigma^2), \quad \mathbf{B} = \mathbf{0},$$

ensuring that the initialization maintains perfect alignment with the pre-trained weights \mathbf{W} . Recent research has focused on developing variants of LoRA aimed at enhancing its performance (Liu et al., 2024; Wang et al., 2024a;b).

3. Method

CLOQ is designed to enhance the fine-tuning of quantized LLMs by incorporating calibration data, building upon OPTQ (Frantar et al., 2022a) with a specific focus on the initialization of LoRA adapters. It utilizes second-order information derived from input activations \mathbf{X} to ensure that the low-rank adapter matrices \mathbf{A} and \mathbf{B} are initialized in a manner that minimizes quantization error, particularly

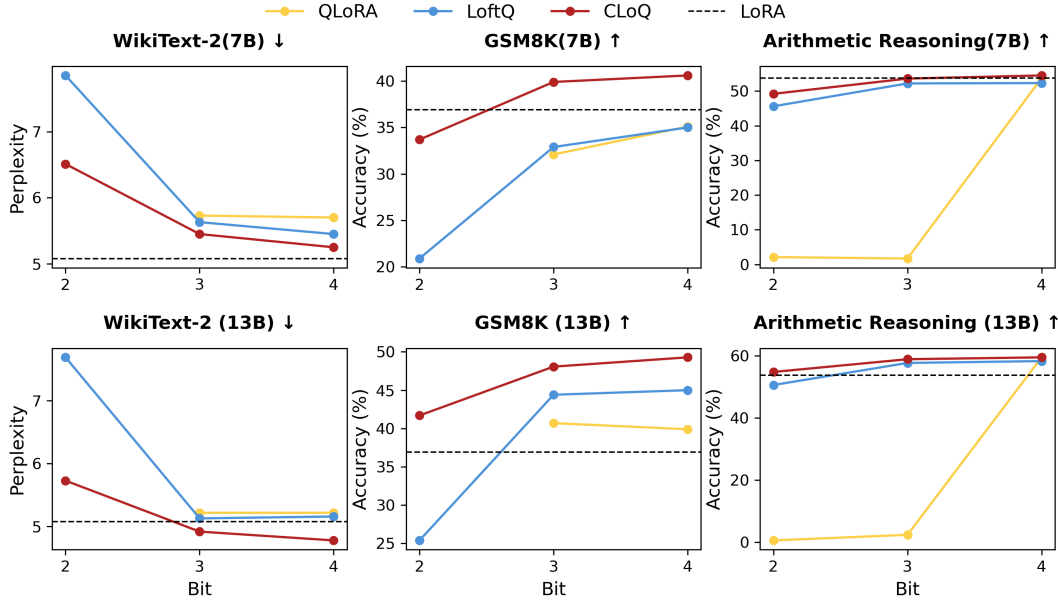


Figure 1. Fine-tuning results of Llama2-7B and Llama2-13B across various tasks. Left: the perplexity measured on WikiText-2. Middle: the accuracy achieved on GSM8K. Right: the average accuracy across multiple arithmetic reasoning tasks.

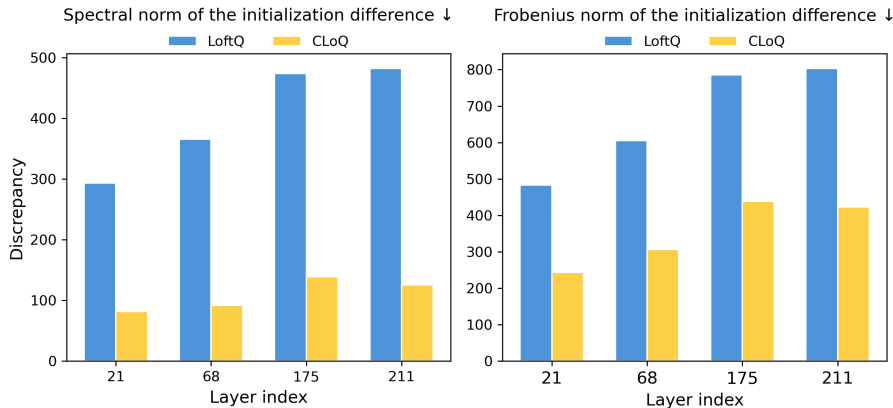


Figure 2. The discrepancy $\|X(Q + AB^T - W)\|$ between the LoRA initialization and the original pre-trained weight matrix, computed using the spectral norm and the Frobenius norm, respectively. The layer shown in the figures above is randomly selected from the Llama2-7B model. The initialization is derived using both CLoQ and LoftQ under INT2 quantization. Notably, CLoQ significantly reduces this discrepancy, demonstrating its effectiveness.

in relation to the data distribution. This alignment allows CLoQ to improve the fine-tuning process by making the low-rank adaptation more responsive to the model’s behavior during calibration, thereby reducing the mismatch between the quantized model and its full-precision counterpart. Importantly, CLoQ uses the same calibration data as OPTQ, ensuring that the initialization of LoRA adapters is both universal and effective across various downstream tasks. In our experiments, we use the Wikitext dataset for calibration, of-

fering a robust and generalizable foundation for fine-tuning across a range of task domains.

3.1. Calibrated Quantization and Low-Rank Initialization

Given the activation matrix X associated with the calibration data, CLoQ aims to solve the following optimization

problem for LoRA initialization:

$$\min_{\mathbf{Q} \in \mathbb{Q}, \mathbf{A} \in \mathbb{R}^{m \times r}, \mathbf{B} \in \mathbb{R}^{n \times r}} \|\mathbf{X}(\mathbf{Q} + \mathbf{A}\mathbf{B}^\top - \mathbf{W})\|_{\mathbb{F}}^2, \quad (2)$$

where $\mathbf{X} \in \mathbb{R}^{(b \cdot l) \times m}$ is the activation matrix associated with a calibration data set of b samples, each represented as an $l \times m$ sub-matrix and stacked along the row dimension. $\mathbb{Q} \subset \mathbb{R}^{m \times n}$ is an appropriate set of all feasible quantized weights. The objective of (2) is to ensure that the initialized weights for the LoRA model, $\mathbf{Q} + \mathbf{A}\mathbf{B}^\top$, closely approximate the pre-trained model weights $\mathbf{W} \in \mathbb{R}^{m \times n}$ when applied to the activation matrix.

This problem can, in theory, be solved using an alternating minimization method, whose t -th iteration reads:

$$\begin{aligned} \mathbf{Q}^{t+1} &= \arg \min_{\mathbf{Q} \in \mathbb{Q}} \|\mathbf{X}(\mathbf{Q} + \mathbf{A}^t(\mathbf{B}^t)^\top - \mathbf{W})\|_{\mathbb{F}}^2 \\ \mathbf{A}^{t+1}, \mathbf{B}^{t+1} &= \arg \min_{\mathbf{A}, \mathbf{B}} \|\mathbf{X}(\mathbf{A}\mathbf{B}^\top + \mathbf{Q}^{t+1} - \mathbf{W})\|_{\mathbb{F}}^2 \end{aligned}$$

In practice, we observe that it suffices to just perform a single iteration with the initialization $\mathbf{A}^0(\mathbf{B}^0)^\top = 0$. Therefore, the proposed CLoQ method comprises two steps: a quantization step and a low-rank approximation step under a linear transformation, which we detail in the rest of this section.

After solving the above problem and obtaining $\mathbf{Q}, \mathbf{A}, \mathbf{B}$, we fix the quantized weights \mathbf{Q} and update only the low-rank components \mathbf{A}, \mathbf{B} during the subsequent fine-tuning stage, as per the LoRA approach.

3.1.1. POST-TRAINING QUANTIZATION

With $\mathbf{A}\mathbf{B}^\top$ initialized to zero, the problem of finding \mathbf{Q} simplifies to the standard layer-wise post-training quantization (PTQ):

$$\min_{\mathbf{Q} \in \mathbb{Q}} \|\mathbf{X}(\mathbf{Q} - \mathbf{W})\|_{\mathbb{F}}^2, \quad (3)$$

an optimization problem that has been extensively studied in recent literature (Frantar et al., 2022b; Chee et al., 2023; Xiao et al., 2023; Zhang et al., 2024b). For this, we adopt the widely-used OPTQ method (Frantar et al., 2022a). To further enhance OPTQ’s performance, we incorporate a pre-processing technique called weight magnitude reduction (MagR) (Zhang et al., 2024a), which modifies \mathbf{W} by removing outliers prior to quantization. MagR preprocessing significantly improves OPTQ’s effectiveness in the low-bit regime while introducing minimal additional time beyond the OPTQ process and incurring no computational or memory overhead during inference time.

3.1.2. GENERALIZED LOW-RANK APPROXIMATION

After obtaining \mathbf{Q} in the quantization step, we denote by

$$\Delta\mathbf{W} = \mathbf{W} - \mathbf{Q}$$

the residual of the quantized weights. To determine \mathbf{A} and \mathbf{B} , we solve the following low-rank approximation problem under the linear transformation induced by \mathbf{X} :

$$\min_{\mathbf{A} \in \mathbb{R}^{m \times r}, \mathbf{B} \in \mathbb{R}^{n \times r}} \|\mathbf{X}(\mathbf{A}\mathbf{B}^\top - \Delta\mathbf{W})\|_{\mathbb{F}}^2. \quad (4)$$

It is important to note that problem (4) is non-trivial due to the presence of the matrix \mathbf{X} , and its optimal solution is not given by directly computing the low-rank approximation (or SVD) of $\Delta\mathbf{W}$. However, we demonstrate in the following result that the problem can be solved accurately by performing just two SVDs:

Theorem 3.1. *Suppose the activation matrix $\mathbf{X} \in \mathbb{R}^{(b \cdot l) \times m}$ ($b \cdot l \gg m$) is of full-rank. Suppose the Gram (or Hessian) matrix $\mathbf{H} = \mathbf{X}^\top \mathbf{X} \in \mathbb{R}^{m \times m}$ has the SVD $\mathbf{H} = \mathbf{U}_H \Sigma_H \mathbf{U}_H^\top$ and denote by $\mathbf{R} = \Sigma_H^{\frac{1}{2}} \mathbf{U}_H^\top$ the non-symmetric root of \mathbf{H} . Then any pair (\mathbf{A}, \mathbf{B}) satisfying*

$$\mathbf{A}\mathbf{B}^\top = \mathbf{R}^{-1} \text{LR}_r(\mathbf{R}\Delta\mathbf{W}). \quad (5)$$

permits an optimal solution to problem (4). Here $\text{LR}_r(\mathbf{R}\Delta\mathbf{W}) = \arg \min_{\text{rank}(\mathbf{Z}) \leq r} \|\mathbf{Z} - \mathbf{R}\Delta\mathbf{W}\|_{\mathbb{F}}^2$ denotes the best rank- r approximation of $\mathbf{R}\Delta\mathbf{W}$.

Proof. Firstly, we observe that the objective in (4) has the following equivalent expression:

$$\begin{aligned} & \|\mathbf{X}(\mathbf{A}\mathbf{B}^\top - \Delta\mathbf{W})\|_{\mathbb{F}}^2 \\ &= \text{Tr}((\mathbf{A}\mathbf{B}^\top - \Delta\mathbf{W})^\top \mathbf{X}^\top \mathbf{X} (\mathbf{A}\mathbf{B}^\top - \Delta\mathbf{W})) \\ &= \text{Tr}((\mathbf{A}\mathbf{B}^\top - \Delta\mathbf{W})^\top \mathbf{H} (\mathbf{A}\mathbf{B}^\top - \Delta\mathbf{W})) \\ &= \text{Tr}((\mathbf{A}\mathbf{B}^\top - \Delta\mathbf{W})^\top \mathbf{R}^\top \mathbf{R} (\mathbf{A}\mathbf{B}^\top - \Delta\mathbf{W})) \\ &= \|\mathbf{R}(\mathbf{A}\mathbf{B}^\top - \Delta\mathbf{W})\|_{\mathbb{F}}^2 \\ &= \|\mathbf{R}\mathbf{A}\mathbf{B}^\top - \mathbf{R}\Delta\mathbf{W}\|_{\mathbb{F}}^2, \end{aligned}$$

where in the third equality, we used the identity: $\mathbf{H} = \mathbf{R}^\top \mathbf{R}$. Moreover, since \mathbf{X} is full rank, \mathbf{H} is invertible, and so is \mathbf{R} .

Based on these facts, we interpret problem (4) as finding the standard best rank- r approximation of $\mathbf{R}\Delta\mathbf{W}$, $\text{LR}_r(\mathbf{R}\Delta\mathbf{W})$, which can be obtained by performing the SVD of $\mathbf{R}\Delta\mathbf{W}$ and extracting the top- r principal components. Then, (\mathbf{A}, \mathbf{B}) is an optimal solution to

$$\min_{\mathbf{A} \in \mathbb{R}^{m \times r}, \mathbf{B} \in \mathbb{R}^{n \times r}} \|\mathbf{R}\mathbf{A}\mathbf{B}^\top - \mathbf{R}\Delta\mathbf{W}\|_{\mathbb{F}}^2$$

if and only if

$$\mathbf{R}\mathbf{A}\mathbf{B}^\top = \text{LR}_r(\mathbf{R}\Delta\mathbf{W}).$$

Consequently, since \mathbf{R} is invertible, any (\mathbf{A}, \mathbf{B}) fulfilling

$$\mathbf{A}\mathbf{B}^\top = \mathbf{R}^{-1} \text{LR}_r(\mathbf{R}\Delta\mathbf{W}).$$

permits an optimal solution to problem (4). \square

To apply Theorem 3.1 to solve problem (4), we make the following observations:

- One SVD is required to compute \mathbf{R} and another is needed to determine $\text{LR}_r(\mathbf{R}\Delta\mathbf{W})$. Given that $\mathbf{R} \in \mathbb{R}^{m \times m}$ and $\mathbf{R}\Delta\mathbf{W} \in \mathbb{R}^{m \times n}$, while $\mathbf{X} \in \mathbb{R}^{(b \cdot l) \times m}$, the computational complexity of SVDs required for solving (4) is independent of $b \cdot l$, which is significantly larger than m or n in practice. Here l represents the context length, and b denotes the size of the calibration dataset.
- Indeed, (5) admits infinitely many optimal solutions. Suppose $\text{LR}_r(\mathbf{R}\Delta\mathbf{W})$ has the form $\mathbf{U}_{:r} \Sigma_{:r} \mathbf{V}_{:r}^\top$. In our experiments, we consistently take $\mathbf{A} = \mathbf{R}^{-1} \mathbf{U}_{:r} \Sigma_{:r}$ and $\mathbf{B} = \mathbf{V}_{:r}$, which empirically performs well. However, it is clear that if (\mathbf{A}, \mathbf{B}) is an optimal solution, then for any invertible $\mathbf{C} \in \mathbb{R}^{r \times r}$, the pair $(\mathbf{A}\mathbf{C}, \mathbf{B}(\mathbf{C}^{-1})^\top)$ also satisfies (5) and thus provides an optimal solution to (4). For example, $(\mathbf{A}, \mathbf{B}) = (\mathbf{R}^{-1} \mathbf{U}_{:r}, \mathbf{V}_{:r} \Sigma_{:r})$ or $(\mathbf{R}^{-1} \mathbf{U}_{:r} \Sigma_{:r}^{\frac{1}{2}}, \mathbf{V}_{:r} \Sigma_{:r}^{\frac{1}{2}})$. In Section 5, our ablation study shows that the combination $(\mathbf{R}^{-1} \mathbf{U}_{:r} \Sigma_{:r}, \mathbf{V}_{:r})$ gives the best practical performance during the subsequent LoRA fine-tuning process. A more comprehensive theoretical analysis of different initialization schemes (Hayou et al., 2024) will be addressed in future work.
- When \mathbf{H} is not invertible or poorly conditioned, we propose adding a small constant λ to the diagonal elements. Specifically, we typically set $\lambda = 0.01 \frac{\text{Tr}(\mathbf{H})}{m}$, following a strategy similar to that used in the prior works (Frantar et al., 2022b; Chee et al., 2023). This adjustment has consistently proven effective in mitigating numerical issues and ensuring stability during computations.

To summarize, our proposed CLoQ method is detailed in Algorithm 1.

4. Experiment

In this section, we evaluate the effectiveness of CLoQ on language modeling, arithmetic reasoning, and commonsense reasoning tasks. The fine-tuning through CLoQ consists of two key stages: the initialization step and the fine-tuning step. In the initialization step, we quantize the full-precision weight \mathbf{W} into low-precision weight \mathbf{Q} and find optimal low-rank matrices \mathbf{A} and \mathbf{B} that minimize the residual error. In the finetuning step, the quantized weight matrix \mathbf{Q} is fixed in low-precision, while \mathbf{A} and \mathbf{B} are trained through back-propagation. The detailed hyperparameter settings for all our experiments are presented in the Appendix A.

Algorithm 1 CLoQ for initializing one linear layer

Input: Gram matrix of activations $\mathbf{H} = \mathbf{X}^\top \mathbf{X} \in \mathbb{R}^{m \times m}$, Pre-trained weight matrix $\mathbf{W} \in \mathbb{R}^{m \times n}$, Rank $r \ll \min(m, n)$.

Output: Quantized weight matrix $\mathbf{Q} \in \mathbb{R}^{m \times n}$, Low-rank components $\mathbf{A} \in \mathbb{R}^{m \times r}$, $\mathbf{B} \in \mathbb{R}^{n \times r}$.

- 1: Solve (3) to obtain the quantized weights \mathbf{Q} .
- 2: Compute the residual of quantized weights $\Delta\mathbf{W} = \mathbf{W} - \mathbf{Q} \in \mathbb{R}^{m \times n}$
- 3: Perform SVD: $\mathbf{H} = \mathbf{U}_H \Sigma_H \mathbf{U}_H^\top$
- 4: Evaluate: $\mathbf{R} = \Sigma_H^{\frac{1}{2}} \mathbf{U}_H^\top$
- 5: Perform the SVD of $\mathbf{R}\Delta\mathbf{W}$ to find its best rank- r approximation: $\text{LR}_r(\mathbf{R}\Delta\mathbf{W}) = \mathbf{U}_{:r} \Sigma_{:r} \mathbf{V}_{:r}^\top$
- 6: Compute the low-rank components:

$$\mathbf{A} = \mathbf{R}^{-1} \mathbf{U}_{:r} \Sigma_{:r}$$

$$\mathbf{B} = \mathbf{V}_{:r}$$

Return: $\mathbf{Q}, \mathbf{A}, \mathbf{B}$

Models and Datasets. We test CLoQ on Llama2-7b and Llama2-13b (Touvron et al., 2023) models. Following prior works (Frantar et al., 2022a), we randomly sample 128 instances, each with a context length of 2048 tokens, from the WikiText-2 dataset (Merity et al., 2016) to serve as the calibration set for quantization. Then, we fine-tune and evaluate the models on WikiText-2 for language modeling. For single arithmetic reasoning tasks, we fine-tune and evaluate on the GSM8K (Cobbe et al., 2021). For multi arithmetic reasoning, we fine-tune the models on Math10K (Hu et al., 2023) and then evaluate the test sets of AQuA (Ling et al., 2017), GSM8K, MAWPS (Koncel-Kedziorski et al., 2016) and SVAMP (Patel et al., 2021). For commonsense reasoning tasks, we fine-tune the models on Commonsense170K (Hu et al., 2023) and evaluate on eight representative tasks: BoolQ (Clark et al., 2019), PIQA (Bisk et al., 2020), SIQA (Sap et al., 2019), HellaSwag (Zellers et al., 2019), WinoGrande (Sakaguchi et al., 2021), ARC-e, ARC-c (Clark et al., 2018) and OBQA (Mihaylov et al., 2018).

Baselines. We compare with LoRA (Hu et al., 2021), QLoRA (Dettmers et al., 2023), GPTQ-LoRA (GPT, 2023) and LoftQ (Li et al., 2023). LoRA is often considered as the benchmark for fine-tuning performance. QLoRA incorporates NF-quantization, with its low-rank initialization aligning with the standard LoRA method. GPTQ-LoRA integrates OPTQ for the base weights and fine-tunes the LoRA component while preserving its low-rank initialization as in the standard LoRA, with the quantized weights kept frozen. In contrast, LoftQ and LQ-LoRA carefully initialize the quantized weight and low-rank matrices by solving some optimization problems to minimize the approximation error.

Table 1. Accuracy on four arithmetic reasoning tasks. The LoRA rank r is 64 for all methods.

Method	Bit	Llama2-7B				Avg. \uparrow	Llama2-13B				Avg. \uparrow
		GSM8K	SVAMP	MAWPS	AQuA		GSM8K	SVAMP	MAWPS	AQuA	
LoRA	16	43.6	59.4	85.0	27.0	53.7	55.3	67.7	87.4	24.4	58.7
QLoRA	4	42.7	58.7	87.3	26.4	53.7	54.8	69.4	87.0	26.8	59.5
GPTQ-LoRA	4	43.0	58.4	86.1	24.3	52.9	53.2	67.5	85.3	25.6	57.9
LoftQ	4	41.7	56.0	86.3	25.3	52.3	54.9	66.5	87.7	23.9	58.3
CLoQ	4	43.6	60.3	87.0	27.6	54.6	55.4	66.9	88.7	27.2	59.5
QLoRA	3	1.4	1.4	0.7	3.4	1.7	0.8	2.5	0.3	6.2	2.4
GPTQ-LoRA	3	38.9	55.7	84.9	23.2	50.7	50.6	65.2	88.0	22.6	56.6
LoftQ	3	39.9	56.3	86.3	26.4	52.2	53.9	66.1	87.0	23.6	57.7
CLoQ	3	42.2	58.9	89.1	24.0	53.6	52.5	66.1	89.1	28.0	58.9
QLoRA	2	0.9	1.5	0.8	5.1	2.1	0.5	0.7	0.1	0.9	0.6
GPTQ-LoRA	2	21.7	39.0	76.6	22.1	39.9	31.9	49.6	82.5	23.6	46.9
LoftQ	2	29.5	45.8	83.6	23.2	45.6	37.0	55.9	87.7	21.7	50.6
CLoQ	2	34.7	52.0	86.1	24.1	49.2	44.6	57.6	88.7	28.4	54.8

Table 2. Finetuning results of WikiText and GSM8K on Llama2-7B and Llama2-13B.

Method	Bit	Llama2-7B		Llama2-13B	
		Wiki (ppl \downarrow)	GSM8K (acc \uparrow)	Wiki (ppl \downarrow)	GSM8K (acc \uparrow)
LoRA	16	5.08	36.9	5.12	45.3
QLoRA	4	5.70	35.1	5.22	39.9
LoftQ	4	5.24	35.0	5.16	45.0
CLoQ	4	5.25	40.6	4.78	49.3
QLoRA	3	5.73	32.1	5.22	40.7
LoftQ	3	5.63	32.9	5.13	44.4
CLoQ	3	5.45	39.9	4.92	48.1
QLoRA	2	N.A.	N.A.	N.A.	N.A.
LoftQ	2	7.85	20.9	7.69	25.4
CLoQ	2	6.51	33.7	5.73	41.7

4.1. Implementation details

Quantization. We quantize the weights of all linear layers in the base model using MagR preprocessing (Zhang et al., 2024a) followed by OPTQ (Frantar et al., 2022a). The quantization scheme employs uniform (a.k.a. INT) and asymmetric quantization, with a default group size of 64. After quantization, we compute the LoRA components A and B , which are maintained in FP16 precision.

Fine-tuning. Following prior works (Detmeters et al., 2023; Li et al., 2023; Liao et al., 2024), we fine-tune the models using the standard LoRA configuration, with modifications to the LoRA initialization and learning rates. The quantized weights remain fixed, and only the LoRA adapter matrices are trainable during fine-tuning. The rank of the LoRA adapters is consistently set to 64 across all methods. For optimization, we use AdamW (Loshchilov, 2017). All experiments are conducted on NVIDIA A100 GPUs with 80GB of memory.

4.2. Fine-tuning Results

Language modeling. We evaluate the models by reporting the perplexity of language generation on WikiText-2. As shown in Table 2, CLoQ consistently delivers the best per-

formance across the majority of bit levels when compared to other methods. Notably, at INT2, CLoQ demonstrates its effectiveness by achieving a perplexity improvement of 1.34 on Llama2-7B over LoftQ. This result highlights CLoQ’s ability to maintain superior performance even under ultra-low bit quantization constraints.

Arithmetic reasoning (single task). To evaluate the models on GSM8K, we extract numerical answers from the generated solutions and determine accuracy by analyzing these extracted values. As shown in Table 2, CLoQ achieves better performance across different model sizes and quantization bit levels. CLoQ reaches an accuracy of 33.7% at INT2 on Llama2-7B, which surpasses that of LoRA at INT4 across all model sizes. More remarkably, the accuracy of INT2 CLoQ on the Llama2-13B model outperforms that of INT4 QLoRA. On the Llama2-7B model, the accuracy of INT2 CLoQ is better than that of INT3 QLoRA and is comparable with INT4 QLoRA.

Arithmetic reasoning. To evaluate CLoQ’s effectiveness across multiple arithmetic reasoning tasks, we fine-tuned models on the Math10K dataset and assessed their performance on four separate math reasoning benchmarks, demonstrating their adaptability to diverse mathematical challenges. As shown in Table 1, CLoQ consistently outperforms other methods across various model sizes and quantization levels, achieving superior accuracy both on average and on individual datasets. Notably, at INT2 quantization, CLoQ delivers substantial gains, improving accuracy by 3.6% and 4.8% over LoftQ on Llama2-7B and Llama2-13B, respectively. Even at INT4, CLoQ surpasses both LoRA and QLoRA on Llama2-7B. Moreover, for complex reasoning tasks such as GSM8K and SVAMP, CLoQ achieves 5.2% and 6.2% higher accuracy than LoftQ on Llama2-7B, highlighting its robustness in challenging problem settings. These results underscore CLoQ’s remarkable potential in handling intricate reasoning tasks with enhanced accuracy, even under ultra-low-bit quantization.

Table 3. Accuracy on eight commonsense reasoning tasks. The LoRA rank $r = 64$ for all methods.

Model	Method	Bit	BoolQ	PIQA	SIQA	HellaS.	WinoG.	ARC-e	ARC-c	OBQA	Avg. ↑
Llama2-7B	LoRA	16	73.6	86.5	81.8	95.2	86.9	89.4	76.7	86.7	84.6
	QLoRA	4	73.9	84.4	79.7	93.3	84.6	86.1	73.0	85.1	82.5
	GPTQ-LoRA	4	73.4	83.6	79.3	93.3	84.5	86.5	72.8	83.3	82.1
	LoftQ	4	73.7	86.0	81.1	94.6	86.3	88.1	75.5	86.2	83.9
	CLoQ	4	74.2	86.3	81.6	95.1	85.9	88.7	75.7	86.4	84.2
	GPTQ-LoRA	3	71.8	82.7	79.3	92.1	82.8	84.2	70.6	83.4	80.8
	LoftQ	3	74.0	85.6	81.0	94.3	85.6	88.1	75.4	85.5	83.7
	CLoQ	3	73.5	86.1	81.2	94.8	85.4	88.6	75.1	85.0	83.7
	GPTQ-LoRA	2	62.2	49.5	33.3	25.1	49.4	25.0	22.6	27.6	36.8
	LoftQ	2	62.4	70.5	73.4	78.8	71.0	66.5	50.8	62.3	67.0
	CLoQ	2	70.2	82.2	78.9	91.7	82.2	83.6	70.7	81.4	80.1
	Llama2-13B	LoRA	16	76.3	88.5	83.4	96.5	89.6	92.8	81.7	89.6
QLoRA		4	74.9	86.6	81.5	94.9	86.9	89.1	77.1	87.2	84.8
GPTQ-LoRA		4	74.5	86.1	81.8	94.7	86.8	89.0	77.1	84.5	84.3
LoftQ		4	76.0	87.9	82.8	95.8	88.9	91.2	80.8	88.8	86.5
CLoQ		4	75.7	88.4	82.9	96.3	89.4	91.1	81.9	90.0	87.0
GPTQ-LoRA		3	73.5	85.2	81.1	94.1	85.7	87.9	75.5	85.3	83.5
LoftQ		3	75.2	87.8	82.8	96.3	89.5	91.1	81.4	88.0	86.5
CLoQ		3	75.3	88.1	82.8	95.9	90.1	91.2	80.7	90.6	86.8
GPTQ-LoRA		2	62.2	50.1	34.0	25.1	49.6	25.0	22.7	27.6	37.1
LoftQ		2	65.9	76.4	78.0	84.4	76.1	75.1	60.1	72.7	73.6
CLoQ		2	73.9	85.5	81.6	94.8	87.3	89.5	77.4	84.8	84.4

Table 4. Finetuning results of WikiText and GSM8K for g128 and per-channel quantization.

Method	Bit	Llama2-7B		Llama2-13B	
		Wiki (ppl↓)	GSM8K (acc↑)	Wiki (ppl↓)	GSM8K (acc↑)
LoRA	16	5.08	36.9	5.12	45.3
LoftQ	4	5.24	35.0	5.16	45.0
CLoQ	4	5.28	39.5	4.79	49.3
LoftQ	3	5.63	32.9	5.13	44.4
CLoQ	3	5.55	38.7	4.99	46.9
	3g128	5.48	39.0	4.94	48.6
LoftQ	2	7.85	20.9	7.69	25.4
CLoQ	2	6.95	31.6	6.07	40.0
	2g128	6.77	32.1	5.90	40.4

Commonsense reasoning. We further evaluate the effectiveness of CLoQ on commonsense reasoning tasks by fine-tuning the models on the Commonsense170K dataset and testing their performance across eight reasoning benchmarks, as presented in Table 3. Consistent with its performance on arithmetic reasoning tasks, CLoQ outperforms other methods across different model sizes and quantization levels, achieving notable improvements both on average and within each dataset.

At INT2, CLoQ delivers a substantial performance boost over LoftQ, with an average accuracy improvement exceeding 10%, even approaching the performance levels typically observed with INT4 QLoRA. For instance, INT2 CLoQ exhibits only a minimal average accuracy drop of 0.04% compared to INT4 QLoRA. Additionally, at INT4, CLoQ attains impressive accuracy, reducing the gap to FP16 LoRA to just 0.4%. These results indicate that CLoQ significantly enhances LoRA’s learning capacity, enabling it to better

adapt to a diverse range of tasks.

4.3. Ablation study

Quantization with different group sizes. We examine the performance of CLoQ for group-wise-128 (g128) and per-channel quantization. Table 4 shows the results for Wikitext-2 and GSM8K. Table 5 shows the result on multiple arithmetic reasoning tasks. CLoQ demonstrates superior performance even under per-channel quantization, achieving results remarkably close to those of g64 quantization. Moreover, it consistently outperforms LoftQ in these settings, demonstrating its effectiveness across all quantization bit levels.

LoRA initialization with different (A, B) combinations. Furthermore, we investigate the performance of various combinations of (A, B) in Algorithm 1, as shown in Table 6. We tried four different combinations of (A, B) , including $(R^{-1}U_{:r}, V_{:r}\Sigma_{:r})$, $(R^{-1}U_{:r}\Sigma_{:r}^{\frac{1}{2}}, V_{:r}\Sigma_{:r}^{\frac{1}{2}})$, $(R^{-1}U_{:r}\Sigma_{:r}^{\frac{1}{2}}H_{\text{had}}, V_{:r}\Sigma_{:r}^{\frac{1}{2}}H_{\text{had}})$ with H_{had} being the order- r ($r = 64$) Hadamard matrix, and the default choice $(R^{-1}U_{:r}\Sigma_{:r}, V_{:r})$. Table 6 shows that the default combination of initialized adapters gives the best performance in the subsequent fine-tuning phase.

5. Related Work

When quantizing pre-trained models, QLoRA (Dettmers et al., 2023) primarily emphasizes the quantization process, often overlooking the critical importance of subsequent LoRA fine-tuning. It adopts the fixup initialization strategy

Table 5. Accuracy on four arithmetic reasoning tasks for g128 and per-channel quantization. The LoRA rank is $r = 64$.

Method	Bit	Llama2-7B					Llama2-13B				
		GSM8K	SVAMP	MAWPS	AQuA	Avg. \uparrow	GSM8K	SVAMP	MAWPS	AQuA	Avg. \uparrow
LoRA	16	43.6	59.4	85.0	27.0	53.7	55.3	67.7	87.4	24.4	58.7
CLOQ	4	42.8	59.5	88.2	25.2	53.9	53.7	67.2	89.1	26.8	59.2
CLOQ	3	41.4	58.5	86.6	24.4	52.7	51.4	65.2	88.2	27.6	58.1
	3g128	42.2	58.9	89.1	24.0	53.6	52.8	64.7	89.5	27.2	58.3
CLOQ	2	34.0	49.1	87.0	24.8	48.7	40.5	55.2	88.7	28.0	53.1
	2g128	33.9	46.9	85.7	26.4	48.9	41.4	58.0	89.1	25.2	53.4

Table 6. Fine-tuning results of different combinations of (A, B) on WikiText-2 and GSM8K. The LoRA rank is $r = 64$.

Method	Bit	Llama2-7B	
		Wiki (ppl \downarrow)	GSM8K (acc \uparrow)
LoRA	16	5.08	36.9
$(R^{-1}U_{:r}, V_{:r}\Sigma_{:r})$	2	880.6	1.6
$(R^{-1}U_{:r}\Sigma_{:r}^{\frac{1}{2}}, V_{:r}\Sigma_{:r}^{\frac{1}{2}})$	2	6.68	12.9
$(R^{-1}U_{:r}\Sigma_{:r}^{\frac{1}{2}}H_{\text{had}}, V_{:r}\Sigma_{:r}^{\frac{1}{2}}H_{\text{had}})$	2	6.51	29.0
$(R^{-1}U_{:r}\Sigma_{:r}, V_{:r})$	2	6.51	33.7

used in LoRA, attaching zero-initialized low-rank adapters to the quantized pre-trained model. However, the discrepancies introduced by quantization, especially in extremely low-bit regimes, can significantly impact the initialization of LoRA fine-tuning, ultimately affecting the overall fine-tuning performance. LoftQ (Li et al., 2023) jointly optimizes the quantized weights Q and the low-rank adapter matrices A and B by solving the following optimization problem:

$$\min_{Q, A, B} \|Q + AB^T - W\|_F^2. \tag{6}$$

This approach ensures that Q , A , and B are initialized to minimize the reconstruction error between the quantized and pre-trained weights. By aligning the quantized model’s initial state more closely with its pre-trained counterpart, LoftQ enhances fine-tuning performance without requiring calibration data.

LQ-LoRA (Guo et al., 2024b) assigns importance weights to each parameter by evaluating its sensitivity to output variations, thereby guiding the decomposition process toward critical regions. Specifically, the Fisher matrix is used to weight the reconstruction objective during decomposition. To solve the resulting weighted SVD problem, LQ-LoRA assumes row and column homogeneity in the Fisher matrix, allowing the use of standard SVD techniques at the cost of theoretical precision. However, this approach has limitations. It relies on approximations rather than solving the weighted SVD problem exactly, which may lead to suboptimal results. Furthermore, computing the Fisher matrix requires back-propagation through the pre-trained model, introducing additional computational overhead. Similar to our work, ApiQ (Liao et al., 2024) also employs an

activation-aware initialization strategy utilizing calibration data. However, it relies on two activation matrices—one obtained from the pre-trained model and the other from the quantized model, whereas CLOQ uses only a single pre-trained activation matrix. ApiQ optimizes the discrepancy using standard back-propagation, whereas CLOQ adopts a fully gradient-free approach. By avoiding the computational overhead of back-propagation, CLOQ enables faster adaptation while maintaining high performance across various tasks.

6. Concluding Remarks

In this work, we introduced CLOQ, an efficient and scalable method for fine-tuning quantized LLMs. By leveraging a small calibration dataset, CLOQ optimally initializes LoRA adapters through a novel layer-wise, data-driven approach, significantly improving the fine-tuning process without the need for back-propagation. The use of a closed-form solution for low-rank approximation, computed via two SVDs, ensures that CLOQ is both computationally efficient and highly effective, particularly at ultra-low bit-widths. Our extensive experiments on multiple benchmark datasets demonstrate that CLOQ consistently outperforms existing LoRA-based methods for quantized models, such as QLoRA, in tasks requiring fine-grained precision. The results underscore the potential of CLOQ to enhance the performance of quantized models across a variety of downstream applications, including those that demand high accuracy, like arithmetic reasoning tasks. The simplicity and efficiency of CLOQ make it a promising approach for fine-tuning large-scale quantized LLMs in resource-constrained environments. Future work could further investigate the theoretical implications of different decompositions of the adapter matrices in CLOQ, and how these variations influence the performance of subsequent fine-tuning.

Acknowledgements

This work was partially supported by NSF grants DMS-2208126, DMS-2110836, IIS-2110546, SUNY-IBM AI Research Alliance Grant, and a start-up grant from SUNY Albany. We would also like to thank SUNY Albany for

providing access to the Nvidia A100 GPUs.

References

- Gptqlora: Efficient finetuning of quantized llms with gptq. GitHub repository, 2023. URL <https://github.com/qwopqwop200/gptqlora>. Available at <https://github.com/qwopqwop200/gptqlora>.
- Achiam, J., Adler, S., Agarwal, S., Ahmad, L., Akkaya, I., Aleman, F. L., Almeida, D., Altenschmidt, J., Altman, S., Anadkat, S., et al. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023.
- Behdin, K., Acharya, A., Gupta, A., Keerthi, S., and Mazumder, R. Quantease: Optimization-based quantization for language models—an efficient and intuitive algorithm. *arXiv preprint arXiv:2309.01885*, 2023.
- Bisk, Y., Zellers, R., Gao, J., Choi, Y., et al. Piqa: Reasoning about physical commonsense in natural language. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pp. 7432–7439, 2020.
- Chee, J., Cai, Y., Kuleshov, V., and De Sa, C. M. Quip: 2-bit quantization of large language models with guarantees. In *Advances in Neural Information Processing Systems*, 2023.
- Choi, J., Wang, Z., Venkataramani, S., Chuang, P. I.-J., Srinivasan, V., and Gopalakrishnan, K. Pact: Parameterized clipping activation for quantized neural networks. *arXiv preprint arXiv:1805.06085*, 2018.
- Clark, C., Lee, K., Chang, M.-W., Kwiatkowski, T., Collins, M., and Toutanova, K. Boolq: Exploring the surprising difficulty of natural yes/no questions. *arXiv preprint arXiv:1905.10044*, 2019.
- Clark, P., Cowhey, I., Etzioni, O., Khot, T., Sabharwal, A., Schoenick, C., and Tafjord, O. Think you have solved question answering? try arc, the ai2 reasoning challenge. *arXiv preprint arXiv:1803.05457*, 2018.
- Cobbe, K., Kosaraju, V., Bavarian, M., Chen, M., Jun, H., Kaiser, L., Plappert, M., Tworek, J., Hilton, J., Nakano, R., et al. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*, 2021.
- Dettmers, T., Pagnoni, A., Holtzman, A., and Zettlemoyer, L. Qlora: Efficient finetuning of quantized llms. *arXiv preprint arXiv:2305.14314*, 2023.
- Frantar, E., Ashkboos, S., Hoefler, T., and Alistarh, D. Gptq: Accurate post-training quantization for generative pre-trained transformers. *arXiv preprint arXiv:2210.17323*, 2022a.
- Frantar, E., Ashkboos, S., Hoefler, T., and Alistarh, D. Optq: Accurate quantization for generative pre-trained transformers. In *The Eleventh International Conference on Learning Representations*, 2022b.
- Guo, D., Zhu, Q., Yang, D., Xie, Z., Dong, K., Zhang, W., Chen, G., Bi, X., Wu, Y., Li, Y., et al. Deepseek-coder: When the large language model meets programming—the rise of code intelligence. *arXiv preprint arXiv:2401.14196*, 2024a.
- Guo, H., Greengard, P., Xing, E., and Kim, Y. LQ-LoRA: Low-rank plus quantized matrix decomposition for efficient language model finetuning. In *The Twelfth International Conference on Learning Representations*, 2024b. URL <https://openreview.net/forum?id=xw29VvOMmU>.
- Hayou, S., Ghosh, N., and Yu, B. The impact of initialization on lora finetuning dynamics. *arXiv preprint arXiv:2406.08447*, 2024.
- Houlsby, N., Giurgiu, A., Jastrzebski, S., Morrone, B., De Laroussilhe, Q., Gesmundo, A., Attariyan, M., and Gelly, S. Parameter-efficient transfer learning for nlp. In *International conference on machine learning*, pp. 2790–2799. PMLR, 2019.
- Hu, E. J., Shen, Y., Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., and Chen, W. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*, 2021.
- Hu, Z., Wang, L., Lan, Y., Xu, W., Lim, E.-P., Bing, L., Xu, X., Poria, S., and Lee, R. K.-W. Llm-adapters: An adapter family for parameter-efficient fine-tuning of large language models. *arXiv preprint arXiv:2304.01933*, 2023.
- Hubara, I., Courbariaux, M., Soudry, D., El-Yaniv, R., and Bengio, Y. Quantized neural networks: Training neural networks with low precision weights and activations. *Journal of Machine Learning Research*, 18(187):1–30, 2018.
- Jiang, A. Q., Sablayrolles, A., Mensch, A., Bamford, C., Chaplot, D. S., Casas, D. d. l., Bressand, F., Lengyel, G., Lample, G., Saulnier, L., et al. Mistral 7b. *arXiv preprint arXiv:2310.06825*, 2023.
- Koncel-Kedziorski, R., Roy, S., Amini, A., Kushman, N., and Hajishirzi, H. Mawps: A math word problem repository. In *Proceedings of the 2016 conference of the north american chapter of the association for computational linguistics: human language technologies*, pp. 1152–1157, 2016.

- Li, Y., Yu, Y., Liang, C., He, P., Karampatziakis, N., Chen, W., and Zhao, T. Loftq: Lora-fine-tuning-aware quantization for large language models. *arXiv preprint arXiv:2310.08659*, 2023.
- Liao, B., Herold, C., Khadivi, S., and Monz, C. Apiq: Finetuning of 2-bit quantized large language model. *arXiv preprint arXiv:2402.05147*, 2024.
- Ling, W., Yogatama, D., Dyer, C., and Blunsom, P. Program induction by rationale generation: Learning to solve and explain algebraic word problems. *arXiv preprint arXiv:1705.04146*, 2017.
- Liu, S.-Y., Wang, C.-Y., Yin, H., Molchanov, P., Wang, Y.-C. F., Cheng, K.-T., and Chen, M.-H. Dora: Weight-decomposed low-rank adaptation. *arXiv preprint arXiv:2402.09353*, 2024.
- Loshchilov, I. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017.
- Merity, S., Xiong, C., Bradbury, J., and Socher, R. Pointer sentinel mixture models. *arXiv preprint arXiv:1609.07843*, 2016.
- Mihaylov, T., Clark, P., Khot, T., and Sabharwal, A. Can a suit of armor conduct electricity? a new dataset for open book question answering. *arXiv preprint arXiv:1809.02789*, 2018.
- Patel, A., Bhattamishra, S., and Goyal, N. Are nlp models really able to solve simple math word problems? *arXiv preprint arXiv:2103.07191*, 2021.
- Rastegari, M., Ordonez, V., Redmon, J., and Farhadi, A. Xnor-net: Imagenet classification using binary convolutional neural networks. In *European conference on computer vision*, pp. 525–542. Springer, 2016.
- Sakaguchi, K., Bras, R. L., Bhagavatula, C., and Choi, Y. Winogrande: An adversarial winograd schema challenge at scale. *Communications of the ACM*, 64(9):99–106, 2021.
- Sap, M., Rashkin, H., Chen, D., LeBras, R., and Choi, Y. Socialiqa: Commonsense reasoning about social interactions. *arXiv preprint arXiv:1904.09728*, 2019.
- Shao, W., Chen, M., Zhang, Z., Xu, P., Zhao, L., Li, Z., Zhang, K., Gao, P., Qiao, Y., and Luo, P. Omniquant: Omnidirectionally calibrated quantization for large language models. *arXiv preprint arXiv:2308.13137*, 2023.
- Touvron, H., Martin, L., Stone, K., Albert, P., Almahairi, A., Babaei, Y., Bashlykov, N., Batra, S., Bhargava, P., Bhosale, S., et al. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*, 2023.
- Wang, N., Choi, J., Brand, D., Chen, C.-Y., and Gopalakrishnan, K. Training deep neural networks with 8-bit floating point numbers. *Advances in neural information processing systems*, 31, 2018.
- Wang, S., Yu, L., and Li, J. Lora-ga: Low-rank adaptation with gradient approximation. *arXiv preprint arXiv:2407.05000*, 2024a.
- Wang, Z., Liang, J., He, R., Wang, Z., and Tan, T. Lora-pro: Are low-rank adapters properly optimized? *arXiv preprint arXiv:2407.18242*, 2024b.
- Xiao, G., Lin, J., Seznec, M., Wu, H., Demouth, J., and Han, S. Smoothquant: Accurate and efficient post-training quantization for large language models. In *International Conference on Machine Learning*, pp. 38087–38099. PMLR, 2023.
- Yao, Z., Wu, X., Li, C., Youn, S., and He, Y. Zeroquant-v2: Exploring post-training quantization in llms from comprehensive study to low rank compensation. *arXiv preprint arXiv:2303.08302*, 2023.
- Zellers, R., Holtzman, A., Bisk, Y., Farhadi, A., and Choi, Y. Hellaswag: Can a machine really finish your sentence? *arXiv preprint arXiv:1905.07830*, 2019.
- Zhang, A., Wang, N., Deng, Y., Li, X., Yang, Z., and Yin, P. Magr: Weight magnitude reduction for enhancing post-training quantization. In *Advances in neural information processing systems*, 2024a.
- Zhang, A., Yang, Z., Wang, N., Qi, Y., Xin, J., Li, X., and Yin, P. Comq: A backpropagation-free algorithm for post-training quantization. *arXiv preprint arXiv:2403.07134*, 2024b.

A. Experimental Details

A.1. Language modeling

To study the capability of CLoQ, we fine-tune quantized models on the WikiText-2 training set and measure perplexity on the validation set. The hyper-parameters used for fine-tuning are provided in Table 7 and Table 8. We evaluate the models on the validation set at each epoch and report the lowest achieved perplexity.

A.2. Arithmetic reasoning

Single task (GSM8K)

To assess CLoQ’s arithmetic reasoning capability, we fine-tune quantized models using the GSM8K training set and evaluate their accuracy on the test set. The hyperparameters used for fine-tuning are detailed in Table 7 and Table 8. Model performance is evaluated at each epoch on the test set, and we report the highest recorded accuracy.

Multiple task

Following the framework proposed by (Hu et al., 2023), we adopt a more integrated approach by training a single model across multiple tasks. Specifically, we fine-tune Llama2-7B and Llama2-13B on Math10K, a dataset that aggregates training samples from GSM8K, MAWPS, MAWPS-single, and AQuA. After finetuning, the models are tested on the evaluation sets of AQuA, GSM8K, MAWPS, and SVAMP. The hyper-parameters used for fine-tuning are detailed in Table 7 and Table 8. Moreover, instead of conducting evaluations at every epoch, we assess model performance only after the final epoch.

A.3. Commonsense reasoning

To evaluate the commonsense reasoning capabilities of CLoQ, we consider eight key benchmark tasks: BoolQ, PIQA, SIQA, HellaSwag, WinoGrande, ARC-e, ARC-c, and OBQA. We adopt the framework proposed by (Hu et al., 2023) and fine-tune a single model across all these tasks instead of training separate models. We fine-tune Llama2-7B and Llama2-13B on the merged training set and measure accuracy on the corresponding test sets. The hyper-parameters used for fine-tuning are detailed in Table 7 and Table 8. For evaluation, we forgo per-epoch assessments and instead report the final model’s performance after the last epoch.

Table 7. Hyper-parameter for the finetuning of Llama2.

Hyper-parameter	WikiText-2	GSM8K	Arithmetic reasoning	Commonsense reasoning
Optimizer	AdamW		AdamW	
Weight decay	0.1		1.0	
LR scheduler	cosine		linear	
Warmup ratio	3%		10%	
Epochs	3	6	3	
Batch size	64	32	16	
Max sequence length	1024	512	512	

Table 8. Best learning rate for Llama2-7B and Llama2-13B on the WikiText-2, GSM8K, and multiple Arithmetic Reasoning tasks.

Group size	Task	Llama2-7B			Llama2-13B		
		4 Bits	3 Bits	2 Bits	4 Bits	3 Bits	2 Bits
64	WikiText-2	7e-4	7e-4	6e-4	2e-4	4e-4	4e-4
	GSM8K	3e-4	3e-4	3e-4	4e-4	4e-4	3e-4
	Arithmetic reasoning	5e-4	9e-4	4e-4	2e-4	4e-4	5e-4
	Commonsense reasoning	8e-5	1e-4	4e-5	5e-5	7e-5	5e-5
128	WikiText-2	-	7e-4	5e-4	-	2e-4	5e-4
	GSM8K	-	7e-4	5e-4	-	5e-4	4e-4
	Arithmetic reasoning	-	7e-4	6e-4	-	3e-4	5e-4
per-channel	WikiText-2	7e-4	4e-4	4e-4	2e-4	2e-4	5e-4
	GSM8K	4e-4	4e-4	4e-4	4e-4	5e-4	5e-4
	Arithmetic reasoning	6e-4	8e-4	3e-4	2e-4	4e-4	5e-4