

# PocketLLM: Ultimate Compression of Large Language Models via Meta Networks

Ye Tian<sup>1</sup>, Chengcheng Wang<sup>2</sup>, Jing Han<sup>3</sup>, Yehui Tang<sup>1</sup>, Kai Han<sup>\*1</sup>

<sup>1</sup> Huawei Noah's Ark Lab

<sup>2</sup> University of Sydney

<sup>3</sup> School of Artificial Intelligence, Beijing University of Posts and Telecommunications  
{tianye76, kai.han}@huawei.com, hanj@bupt.edu.cn

## Abstract

As Large Language Models (LLMs) continue to grow in size, storing and transmitting them on edge devices becomes increasingly challenging. Traditional methods like quantization and pruning struggle to achieve extreme compression of LLMs without sacrificing accuracy. In this paper, we introduce PocketLLM, a novel approach to compress LLMs in a latent space via meta-networks. A simple encoder network is proposed to project the weights of LLMs into discrete latent vectors, which are then represented using a compact codebook. A lightweight decoder network is employed to map the codebook's representative vectors back to the original weight space. This method allows for significant compression of the large weights in LLMs, consisting solely of a small decoder, a concise codebook, and an index. Extensive experiments show that PocketLLM achieves superior performance even at significantly high compression ratios, e.g., compressing Llama 2-7B by 10× with a negligible drop in accuracy.

## Introduction

The rapid development of large language models (LLMs) (OpenAI 2023; Touvron et al. 2023a; Zeng et al. 2023) has been one of the most impressive advancements in artificial intelligence in recent years. These models have grown larger and larger in scale, with their capacity and capabilities advancing at an unprecedented rate, following the scaling law of pretraining (Kaplan et al. 2020). As a result, they have demonstrated remarkable proficiency across a wide range of tasks, including natural language understanding and generation to specialized applications in fields such as personal assistant, finance, education, and so on (OpenAI 2023; Yang, Liu, and Wang 2023; Li et al. 2025).

The increasing scale of these models has not only pushed the boundaries of what is possible but has also raised significant challenges related to efficiency and practical deployment. More and more mobile devices tend to integrate the capabilities of LLMs, such as laptops, smartphones, and autonomous vehicles (Gunter et al. 2024; Xu et al. 2024). However, these devices are often limited in terms of storage space, and the large size of large models makes them

impractical for direct deployment. Additionally, the transmission and regular updating of these models from cloud to devices require substantial network bandwidth, which may not always be available or cost-effective, especially in regions with limited connectivity, seriously affecting the user experience. Consequently, there is a pressing need for extreme compression techniques that can significantly reduce the model's size while preserving its performance, enabling efficient storage, faster transmission, and seamless deployment on resource-constrained devices.

Recent studies have shown that traditional methods including pruning and quantization can preserve the performance of LLMs at a low compression ratio (Kurtic et al. 2023; Wei et al. 2024; Dettmers et al. 2023; Guo et al. 2024; Liao et al. 2025; Huang, Huang, and Wen 2025). However, as the compression ratio continues to increase, there will be a significant decrease in accuracy. As the compression ratio increases, pruning and low-bitwidth quantization methods unavoidably lose some crucial information, thereby affecting accuracy. Meanwhile, these methods often require a large amount of computing expenditure to retrain or fine-tune the LLMs. Though some post-training methods are proposed to compress LLMs without retraining, these methods mostly quantify LLMs with 3-4 bitwidth (Frantar et al. 2022; Shao et al. 2023; Lin et al. 2024; Chee et al. 2023; Kim et al. 2023). With the proposal of Low-Rank Adaptation (LoRA) (Hu et al. 2022), some methods introduce additional trainable parameters after compression to ensure accuracy (Ji et al. 2024; Xu et al. 2023; Wang et al. 2025; Ashkboos et al. 2024). Nevertheless, the accuracy is still limited at extreme compression ratios (e.g.,  $> 10\times$ ), and the complex fine-tuning process also complicates the pipelines.

In this paper, we propose to compress the large language models in latent spaces, instead of directly quantizing or pruning the weight matrix as it is. In particular, we first view the weight matrix as a collection of multiple discrete short vectors given that it is challenging to perform quantization or pruning in units of the entire weight matrix. Considering that the complex relationship and redundancy of weight vectors are difficult to learn in the original linear vector space, we construct a meta encoder network to project the weight vectors into latent vectors. We learn a small codebook to represent these latent vectors discretely. We also construct a meta decoder network to map the discrete vectors back to

\*Corresponding author

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

the original linear space for the weight reconstruction. In this way, we only need to store the meta decoder networks, the codebook, and the corresponding indices of latent vectors, so as to compress the weights of LLMs extremely at the cost of only a few additional parameters. Extensive experiments across various popular benchmarks show that our method achieves new state-of-the-art results at extreme compression ratios. Even at  $16\times$  compression, the performance loss of PocketLLM with fine-tuning is still acceptable.

## Related Works

**Traditional Pruning Methods** Traditional pruning methods fall broadly into two categories, i.e. structured pruning and unstructured pruning. For structured pruning methods, researchers often remove channels or layers with non-critical information. SoBP (Wei et al. 2024) selects the pruning structures based on global first-order information and leverages a local greedy approach to refine them. To mitigate information loss, they adopt module-wise reconstruction as an alternative to fine-tuning. LLM-Pruner (Ma, Fang, and Wang 2023) selectively removes non-critical coupled structures based on gradient information and utilizes Low-Rank Adaptation(LoRA) fine-tuning to recover the performance of pruned models. For unstructured pruning, they zero out the weights of multiple neurons without modifying the model structure. SparseGPT (Frantar and Alistarh 2023) focuses on the extremely large-scale instances of sparse regression and proposes an efficient one-shot pruning method. Wanda (Sun et al. 2023) evaluates the priority of each weight by the product of its magnitude and the norm of the corresponding input feature. For each output of linear layers, the weights with lower priority are pruned.

**Traditional Quantization Methods** Traditional quantization methods can be roughly divided into two categories, i.e. quantization during training and post-training methods. Due to the fact that the former methods require large computational overhead to retrain LLMs during the quantification process, academia focuses more on the study of the latter. GPTQ (Frantar et al. 2022) proposes an efficient post-training method that minimizes layer-wise quantitative loss based on approximate second-order information. ZeroQuant (Yao et al. 2022) proposes a fine-grained hardware-friendly quantization scheme that makes the group-wise quantization for weights and token-wise quantization for activation, respectively. OmniQuant (Shao et al. 2023) proposes an efficient quantization framework that optimizes both weight-only and weight-activation quantization through block-wise error minimization.

**Codebook-based Methods** Codebook-based methods also become a research hotspot in recent years (Tseng et al. 2024a; Liu et al. 2024; Tseng et al. 2024b; Baalen et al. 2024; Egiazarian et al. 2024). The concept of codebook is usually applied in the field of image generation. To compress a database of image vectors or embeddings, researchers learn a relatively small set of discrete vectors called "codebook" and optimize these vectors using nearest neighbor algorithms. For each given database vector, we can choose the nearest vector in the codebook for representation. This compression process is also known as vector quantization (VQ).

For example, AQLM (Egiazarian et al. 2024) constructs an end-to-end quantization method based on multi-group codebooks. It combines vectors from multiple codebooks to represent a weight vector, thereby reducing the difficulty of approximation. VPTQ (Liu et al. 2024) utilizes Second Order Optimization with brief codebooks to refine the post-training quantization for LLM. QTIP (Tseng et al. 2024b) instead uses trellis coded quantization to separate the codebook size from the bitrate and dimension to improve the performance further. Although compared to traditional methods, the compression rate is significantly improved while the accuracy remains better. However, these methods need to construct complex fine-tuning pipelines to restore the accuracy of the compressed models. Essentially, it is because of the high difficulty of constructing codebooks with good representational ability in the current linear space.

## Approach

### Prelimineries

**Transformer** Transformer is composed of interleavedly stacked multi-head self-attention modules and feed-forward network blocks. The self-attention performs as follows:

$$O = W_o \text{Softmax}(QK^T / \sqrt{d})V, \quad (1)$$

where  $W_o$  is the output projection matrix,  $d$  is the feature dimension of queries,  $Q$ ,  $K$  and  $V$  are obtained by fully-connected layers with weights  $W_q$ ,  $W_k$  and  $W_v$  from the input  $X$  respectively:

$$Q = W_q X, \quad (2)$$

$$K = W_k X, \quad (3)$$

$$V = W_v X. \quad (4)$$

The number of parameters in the self-attention module is about  $4d^2$ .

Most parameters of the feed-forward network are occupied by the two fully-connected layers:

$$X' = W_{down} \sigma(W_{up} X), \quad (5)$$

where  $W_{up}$  and  $W_{down}$  are the weight matrices, and  $\sigma$  is the nonlinear activation function. The hidden dimension is usually set larger than the input dimension, e.g.,  $4d$ . The number of parameters in the FFN module is about  $8d^2$ . Some variants of the standard FFN have more fully-connected layers like Llama with  $12d^2$  parameters.

### PocketLLM Framework

As Transformers in LLMs continue to scale, storing and transmitting them on edge devices becomes increasingly challenging. Traditional methods such as quantization and pruning often fail to achieve extreme compression without sacrificing accuracy. In this paper, we introduce PocketLLM for compressing LLMs in latent spaces using meta-networks. The entire pipeline is illustrated in Figure 1. First, we divide the weight matrix into multiple vectors and use an encoder network to project these weight vectors into the latent spaces. By gradually incorporating nonlinear factors, the hidden features in the weight vectors are effectively

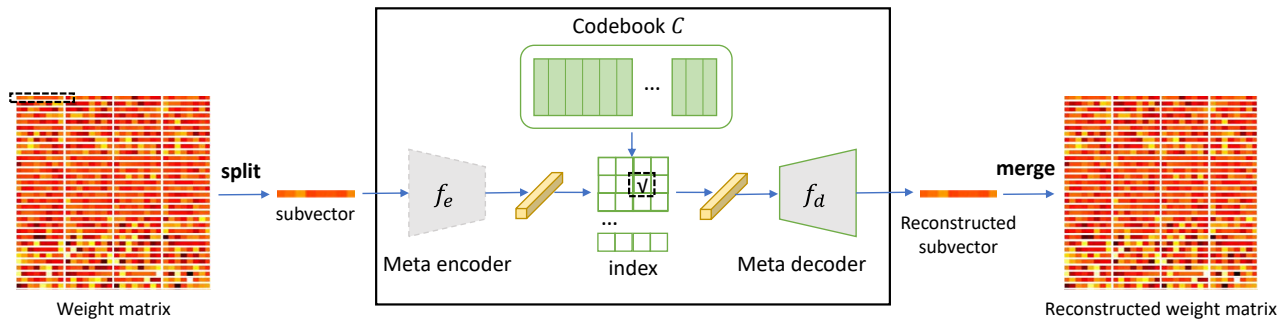


Figure 1: PocketLLM framework: compressing weight vectors in latent spaces with meta-networks.

learned, similar to the process of noise addition in a diffusion model. Next, we build a codebook of discrete latent vectors to represent the weight vectors, which corresponds to a clustering operation in latent spaces, with the number of clusters matching the codebook size. Each latent weight vector is then replaced by its closest representation in the codebook. Finally, a decoder is constructed to map the representation vectors back to the original space, aiming to recover the original model weights. The following sections will provide a detailed explanation of the design of these modules.

### Encoding Weights into Latent Spaces

The majority of weights in large language models come from the mapping matrices of the linear layers, so we should focus on how to compress these weights. We divide the fitting problem of the entire weight matrix into fitting problems of multiple subvectors. Given a weight matrix  $W \in \mathbb{R}^{d_{in} \times d_{out}}$ , we divide each row  $W_i \in \mathbb{R}^{d_{out}}$  into several subvectors so we can operate at a finer granularity:

$$W_i = [W_i^1, W_i^2, \dots, W_i^L], \quad (6)$$

where  $W_i^L \in \mathbb{R}^d$  is the subvector,  $L$  is the number of subvectors and  $d = d_{out}/L$ . For all the weight matrices in a LLM, we can obtain a large number of subvectors, denoted as  $S \in \mathbb{R}^{N \times d}$  where  $N = d_{in}L$  is the total number of subvectors.

For these weight vectors  $S \in \mathbb{R}^{N \times d}$ , each of them can be represented as a nonlinear combination of several intrinsic codewords. The nonlinear relationship is hard to explicitly represent, so we propose learning this relationship in a compact latent space. For each subvector  $S_i$ , a meta encoder  $f_e()$  is utilized to transform the subvector into an embedding in latent spaces:

$$Z_i = f_e(S_i), \quad (7)$$

where  $Z_i \in \mathbb{R}^d$ . The meta encoder  $f_e()$  can be implemented with a neural network such as Multi-layer Perceptron (MLP).

Due to the particularity of using weight subvectors as input, we design a simple yet effective norm network to replace the frequently used Layer Normalization (LN) network, called Reshaped Layer Normalization (RLN). As with previous pre-norm works, we perform a norm operation before residual linking to prevent issues such as gradient explosion. After layer-by-layer nonlinear mapping in the encoder network, we will ultimately project the weight vectors

into latent spaces for representation in the codebook. It is worth noting that the encoder can be discarded after the end of training. Because we only need to retain the representation vectors in the codebook, the index of each weight vector corresponding to the representation vector, and the decoder.

**Reshaped Layer Normalization** Layer Normalization (LN) is usually used in the Transformer networks to normalize the sequence input embeddings. So it is an intuitive idea to utilize LN networks to process the weight subvectors that are also sequence input. After thorough experimentation, we found that this approach effectively avoids problems such as gradient explosion. However, it also causes a significant difference between the weight vector represented by the codebook and the original vector. Considering that for the input  $1 \times d$  sequence (i.e., weight subvector), LN only considers the relationships between  $d$  elements in the sequence and ensures that the distribution of these elements is uniform, without considering elements from other input sequences. However, each  $1 \times d$  vector is only a part artificially split from the row vectors of the weight matrix, and the elements often do not satisfy a certain distribution relationship or have incomplete semantic associations with each other. Because these elements tend to satisfy a certain distribution within the entire row vectors, we believe that normalizing the original row vectors is more meaningful. Based on the above considerations, we reshape the  $1 \times d$  weight vectors back to their original size (i.e. the entire weight row vector) before normalizing. It is equivalent to aligning the vector elements once at the semantic level. After normalization, we then divide the entire row vector into the  $1 \times d$  weight subvectors. We name this operation as Reshaped Layer Normalization (RLN). Extensive experiments show that RLN can greatly improve the effect without introducing additional parameter quantities or computational overhead.

### Compression in Latent Spaces

As mentioned above, a codebook is composed of a set of discrete vector sets in latent spaces. We actually cluster weight vectors in latent spaces to obtain  $K$  category centers. Each weight vector in latent spaces is represented by the class center closest to it, thus achieving the purpose of compression. Specifically, given weight vectors of a certain layer  $Z \in \mathbb{R}^{N \times d}$  in latent spaces. At the same time, we initialize  $1 \times d$  vectors with the quantity of  $K$  in the specific codebook

$C \in \mathbb{R}^{K \times d}$ . For the weight vector  $Z_i$  and codebook  $C$ , we construct a nearest neighbor mapping. Each latent embedding  $Z_i$  is approximated by a codeword (Eq. 8).

$$q(Z'_i = C_k) = \begin{cases} 1, & \text{if } k = \arg \min_j \|Z_i - C_j\|_2 \\ 0, & \text{otherwise,} \end{cases} \quad (8)$$

where  $Z'_i$  is the approximate discrete latent vector of  $Z_i$  based on the codebook. This mapping can be represented by an index array  $I \in \mathbb{R}^N$ , which contains  $N$  integers. In this way, the weight matrix of a certain layer can be compressed as a codebook and an index array.

The forward process in Eq. 8 is non-differential, meaning it cannot be optimized by a standard back-propagation algorithm. We approximate the gradient using straight-through estimator (Bengio, Léonard, and Courville 2013) and directly pass-through the gradient of  $Z'$  to  $Z$ :

$$\frac{\partial \ell}{Z} = \frac{\partial \ell}{Z'}. \quad (9)$$

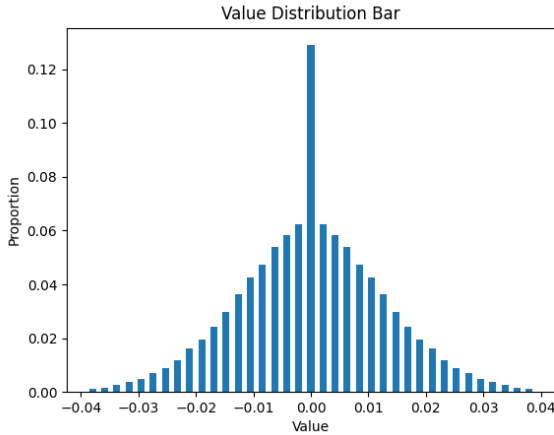


Figure 2: We visualize the numerical distribution of values within the 99.9% range in the value projection weights  $W_v$ .

Meanwhile, we conduct experimental statistical analysis of the numerical distribution in the weight vectors and find that, except for a few outliers, most values closely approximate a normal distribution. As shown in Figure 2, taking a projection layer of MLPs as an example, most values are concentrated in a certain interval and approximately follow a normal distribution. In order to approximate the true spatial distribution of vectors and reduce the difficulty of fitting, we use the normal distribution during initialization. We optimize the clustering algorithm using the simplest nearest neighbor algorithm, and only consider the distance between each weight vector and the cluster center during the training process, decoupling it from the subsequent decoding process. Therefore, we utilize Mean Squared Error (MSE) to optimize the codebook by minimizing the Euclidean distance between each weight vector and its representation vector. Specifically, we have  $Z'_i$  (Eq. 8) as the representative vector for the  $i$ -th latent weight vector  $Z_i$ , the loss function

can be calculated as follows:

$$MSE = \sum_{i=1}^N \|Z_i - Z'_i\|^2. \quad (10)$$

### Weight Reconstruction with Meta Decoders

As we supposed above, the subvector is a complex nonlinear transformation of an intrinsic codeword. This transformation is difficult to be captured with simple operations like linear fitting, so we propose using a meta decoder to reconstruct the original subvector. The approximated codeword is fed into the meta decoder  $f_d()$  for reconstruction:

$$\hat{S}_i = f_d(Z'_i), \quad (11)$$

where  $\hat{S}_i \in \mathbb{R}^d$  is the reconstructed subvector.

In order to maintain consistency with the structure of the encoder, we samely construct a  $m$ -layer MLP and use residual links in every layer except for the last layer. Similarly, we use RLN instead of LN before the residual links. The input of the decoder is the representation vector of each weight vector in the latent space. The decoding process is the process of mapping the representation vectors back to the original vector space and expecting to obtain the same output as the input  $S$  of the Encoder. We denote the output of the decoder as  $\hat{S} \in \mathbb{R}^{N \times d}$  and reduce the difference between  $S$  and  $\hat{S}$  based on Root Mean Squared Error (RMSE). The specific loss function can be calculated as:

$$RMSE = \sqrt{\sum_{i=1}^N \|S_i - \hat{S}_i\|^2}. \quad (12)$$

### Compression Ratio Analysis

This section analyzes the compression ratio of the proposed method. The number of parameters of the original weight matrices is  $Nd$ . For the compressed version, we only need to store the codebook, the indices, and the meta decoder. Their numbers of parameters are  $Kd$ ,  $N$  and  $N_{fd}$ , respectively. Thus, the compression ratio can be calculated as:

$$r = \frac{Nd}{Kd + N + N_{fd}}. \quad (13)$$

Moreover, if the length of the codebook is  $K$ , we can store the indices as integers with  $\log_2(K)$  bits. We perform half-precision quantization on the codebook. Compared to the original FP32 model, we can calculate the compression ratio as follows:

$$r = \frac{32Nd}{16Kd + \log_2(K)N + 32N_{fd}}. \quad (14)$$

Using the Llama 2-7B as the original model for compression, and taking a 3-layer MLP as an example of the meta decoder, its parameter count is 768. We maintain a codebook with  $K = 2^{15}$  and  $d = 8$ . We have the compression ratio for one  $up$  layer of the FFN module:

$$r = \frac{32 \cdot 45.1M}{16 \cdot 2^{15} \cdot 8 + \log_2(2^{15}) \cdot 5.6M + 32 \cdot 768} = 16.4 \quad (15)$$

## Algorithm Process Overview

In this section, we use pseudocode to briefly describe our algorithm process.

---

### Algorithm 1: PocketLLM: Weight Compression for LLMs

---

```
Require: model, calibration_data, k, d, h
1: for  $i = 1, \dots, \text{model.num\_layers}$  do
2:    $f_e, f_d := \text{initialize}(h)$ 
3:    $C := \text{initialize}(k)$ 
4:    $\text{block} := \text{model.get\_block}(i)$ 
5:   for  $\text{layer} \in \text{linear\_layers}(\text{block})$  do
6:      $W := \text{layer.weight}$ 
7:      $S := \text{split}(W, d)$ 
8:      $Z := f_e(S)$ 
9:      $Z' := \text{K-means}(Z, C)$ 
10:     $\hat{S} := f_d(Z')$ 
11:     $L := \text{RMSE}(S, \hat{S}) + \lambda \text{MSE}(Z, Z')$ 
12:    /* After minimizing the loss  $L$  */
13:     $\hat{W} := \text{merge}(\hat{S}, d)$ 
14:     $\text{layer.weight} := \hat{W}$ 
15:   end for
16:   /* save for fine-tuning if necessary */
17:    $\text{model} := \text{LoRA}(\text{model}, \text{calibration\_data})$ 
18: end for
```

---

## Experiments

We first introduce the implementation details. Then we compare PocketLLM against other approaches under both fine-tuning and non-fine-tuning settings. Next, we show our visual results. Finally, we make a detailed ablation analysis.

### Implement Details

We apply the PocketLLM algorithm based on the Llama 2-7B and Qwen 3-14B model provided by Huggingface (Touvron et al. 2023b; Yang et al. 2025). We use GELU activation function (Hendrycks and Gimpel 2016) in MLP to increase the nonlinearity of the network, which helps to perceive potential differences between vectors. We construct a 3-layer MLP network for both encoder and decoder, and each layer except for the first layer uses residual links to improve performance. We set the value of  $k$  (the size of the codebook) and  $d$  (the length of weight vector) to  $(d, k) = (4, 2^{15})$ ,  $(d, k) = (4, 2^{12})$ ,  $(d, k) = (8, 2^{15})$ ,  $(d, k) = (8, 2^{12})$ , corresponding to  $8\times$ ,  $10\times$ ,  $16\times$ , and  $20\times$  compression experiments. To fairly compare with recent methods, we conducted experiments with and without fine-tuning simultaneously. Different from some methods using block-wise or layer-wise fine-tuning step-by-step during the compression (Egiazarian et al. 2024). We only use the standard LoRA algorithm for fine-tuning once a time after compression. Specifically, we set rank and alpha as 32 and 64, respectively. Other parameters are set to the default value. We use the same calibration datasets (redpajama (Weber et al. 2024), alpaca (Taori et al. 2023)) as recent methods for fine-tuning (Deng et al. 2024; Egiazarian et al. 2024).

### Comparisons with State-of-the-Art Methods

In this section, we compare the PocketLLM with other methods on 7 benchmarks. Specifically, we report the accuracy

on five zero-shot tasks, i.e., WinoGrande (Sakaguchi et al. 2021), PiQA (Tata and Patel 2003), HellaSwag (Zellers et al. 2019), ARC-easy (Arc-E), and ARC-challenge (Arc-C) (Boratto et al. 2018), and report the perplexity on WikiText-2 (Merity et al. 2016) and C4 (Raffel et al. 2020) datasets. As with the previous work, the calculation of the *average bits* only takes quantized weights into account.

In the Table 1, we compare the accuracy of PocketLLM and other methods on five zero-shot tasks. According to the results, we can find that PocketLLM achieves better results regardless of fine-tuning or not at different compression ratios. Especially under the  $8\times$  (compared to the float32 model) compression settings, PocketLLM outperforms other methods by a large margin. Even compared to the uncompressed model, PocketLLM without fine-tuning still has a comparable performance. Meanwhile, under the  $10\times$  compression settings, PocketLLM without fine-tuning has a certain loss in accuracy. However, with just a simple fine-tuning, the performance can be well restored. Under the  $16\times$  compression settings, the performance loss of PocketLLM with fine-tuning is still acceptable. When the compression ratio reaches  $20\times$ , PocketLLM is still able to greatly restore the performance through fine-tuning.

In the Table 2, we also extend PocketLLM on the base model Qwen 3-14B (Yang et al. 2025) for  $8\times$  and  $10\times$  compression ratios, respectively. According to the results, we can find that the precision advantage of PocketLLM is significant. Even if compressed  $10\times$  times, the loss of accuracy is still negligible. More compressing results about Llama 3-8B and Llama 1-7B can be found in supplementary materials.

In the Table 3, we also compare the perplexity of PocketLLM and other methods. According to the results, we can find that PocketLLM achieves comparable performance to AQLM, QTIP, and surpass all the other methods. As for the accuracy gap, we analyze that it is due to our insufficient calibration fine-tuning. The rougher precision recovery approach may not be suitable for the perplexity metrics.

### Ablation Study

In this section, we mainly make ablation studies based on quantized Llama 2-7B without fine-tuning under  $8\times$  compression. We first analyze the difficulty of compressing for different types of layers. Then, we analyze the impact of MLP quantity on the reconstruction loss. Next, we analyze the impact of codebook size. Finally, the effectiveness of RLN and codebook initialization is analyzed.

In the Table 4, in addition to the standard 7 layers, we also investigate the losses caused by compressing all the attention layers ( $q, k, v, o$ ) and all the FFN layers ( $gate, up, down$ ), respectively. Unlike previous experiments, we also add the MMLU (Hendrycks et al. 2020) dataset to observe the loss of different layers in difficult scenarios. It can effectively prevent fluctuations on simple datasets from affecting conclusions. According to the results, we can observe interesting phenomena. Firstly, although a single FFN layer has more parameters, its impact on the loss is not much higher than that of a single attention layer. Second, compressing all the attention layers will bring apparent accuracy loss that is almost equivalent to compressing all the FFN layers. There-

Compress Ratio	Method	Avg_bits	WinoGrande $\uparrow$	PiQA $\uparrow$	HellaSwag $\uparrow$	ArcE $\uparrow$	ArcC $\uparrow$	Avg_acc $\uparrow$
-	Llama 2-7B	32	67.25	78.45	56.69	76.01	43.03	64.29
$2\times\sim 8\times$	LLM-Pruner	11.20	64.25	76.01	-	66.62	37.20	-
	Bolaco	11.20	65.67	75.19	-	69.15	38.74	-
	GPTQ	4.00	68.19	76.61	55.44	66.20	36.77	60.64
	SpQR	3.98	66.93	78.35	<u>56.10</u>	69.11	39.68	62.17
	DuQuant	4.00	63.93	75.68	-	50.00	37.46	-
	QuIP#	4.02	66.85	77.91	55.78	68.06	39.68	61.66
	AQLM	4.04	67.32	78.24	55.99	<u>70.16</u>	<u>41.04</u>	<u>62.55</u>
	VPTQ	4.01	67.10	78.10	55.00	69.00	39.70	61.98
	QTIP	4.01	67.10	<u>78.40</u>	-	68.90	40.00	-
	PocketLLM (ours)	3.98	<b>69.39</b>	<b>78.54</b>	<b>57.45</b>	<b>76.18</b>	<b>43.17</b>	<b>64.95</b>
	LLM-Pruner*	11.20	58.72	71.93	-	61.41	33.96	-
	SoLA*	11.20	62.83	69.97	-	59.13	32.68	-
	Bolaco*	11.20	64.80	71.44	-	66.46	34.39	-
	SoBP*	11.20	66.22	73.50	-	59.81	37.63	-
	Wanda*	8.00	<u>66.61</u>	-	<u>46.45</u>	<u>69.15</u>	35.84	-
	DuQuant*	4.00	62.12	<u>75.24</u>	-	51.89	36.77	-
PocketLLM* (ours)	3.98	<b>68.19</b>	<b>76.93</b>	<b>55.01</b>	<b>73.99</b>	<b>41.38</b>	<b>63.10</b>	
$\sim 10\times$	GPTQ	3.00	59.19	71.49	45.21	58.46	31.06	53.08
	SpQR	2.98	63.54	74.81	51.85	67.42	37.71	59.07
	AQLM	3.04	66.93	76.88	54.12	68.06	38.40	60.88
	GPTVQ	3.13	67.64	77.64	-	<u>72.73</u>	<b>43.69</b>	-
	VPTQ	3.01	<b>68.00</b>	77.30	<u>56.00</u>	69.10	39.30	<u>61.72</u>
	QTIP	3.01	66.90	<u>78.10</u>	-	68.10	38.90	-
	PocketLLM (ours)	2.98	<u>67.40</u>	<b>78.13</b>	<b>57.17</b>	<b>74.12</b>	<u>43.52</u>	<b>64.07</b>
	PocketLLM* (ours)	2.98	<b>63.54</b>	<b>75.41</b>	<b>51.67</b>	<b>69.99</b>	<b>37.03</b>	<b>59.53</b>
$\sim 16\times$	QuIP#	2.02	62.43	71.38	42.94	55.56	28.84	52.23
	AQLM	2.02	<u>65.67</u>	74.76	49.55	63.68	32.76	57.28
	GPTVQ	2.13	60.93	70.73	-	58.08	31.48	-
	VPTQ	2.02	64.33	75.19	<u>52.08</u>	63.80	35.24	<u>58.13</u>
	QTIP	2.00	64.70	<u>75.90</u>	-	<u>67.25</u>	<u>35.70</u>	-
	PocketLLM (ours)	2.02	<b>67.25</b>	<b>76.71</b>	<b>53.24</b>	<b>69.07</b>	<b>36.77</b>	<b>60.61</b>
	PocketLLM* (ours)	2.02	<b>58.72</b>	<b>70.95</b>	<b>44.02</b>	<b>58.16</b>	<b>30.38</b>	<b>52.45</b>
$\sim 20\times$	PocketLLM (ours)	1.60	<b>60.30</b>	<b>73.50</b>	<b>47.63</b>	<b>63.58</b>	<b>32.25</b>	<b>55.45</b>
	PocketLLM* (ours)	1.60	<b>52.01</b>	<b>54.08</b>	<b>27.20</b>	<b>33.04</b>	<b>23.21</b>	<b>37.91</b>

Table 1. Evaluation of compressed Llama 2-7B models for the accuracy of 5 zero-shot tasks under different compression ratios. The Avg\_acc column is the mean of these accuracies. \* means no fine-tuning in the method.    represents the second-best result.

Compress Ratio	Method	Avg_bits	WinoGrande $\uparrow$	PiQA $\uparrow$	HellaSwag $\uparrow$	ArcE $\uparrow$	ArcC $\uparrow$	Avg_acc $\uparrow$
-	Qwen 3-14B	32	74.35	80.58	61.86	83.54	55.80	71.23
$\sim 8\times$	SmoothQuant	4.00	63.20	72.40	46.30	69.60	43.40	58.98
	RTN	4.00	65.40	76.30	53.50	76.40	51.00	64.52
	AWQ	4.00	71.70	<u>79.90</u>	<u>60.60</u>	<u>82.20</u>	<u>54.20</u>	69.72
	GPTQ	4.00	<u>73.20</u>	79.70	<u>60.60</u>	81.50	54.00	69.80
	PocketLLM (ours)	3.98	<b>74.90</b>	<b>80.14</b>	<b>60.64</b>	<b>84.85</b>	<b>55.97</b>	<b>71.30</b>
$\sim 10\times$	AWQ	3.00	<u>62.70</u>	76.00	55.90	73.00	44.60	62.44
	PocketLLM (ours)	3.02	<b>74.20</b>	<b>79.49</b>	<b>57.72</b>	<b>84.22</b>	<b>55.25</b>	<b>70.19</b>

Table 2. Evaluation of compressed Qwen 3-14B models for the accuracy on 5 zero-shot tasks.

fore, we can conclude that even if the number of parameters occupied by attention layers is less than one-third, it is still important. Perhaps this can inspire us to design LLMs with limited parameters (smaller than 2B), and we may be able to lose some parameters in FFN modules to ensure the parameters in attention layers.

In the Table 5, we explore the optimal number of model layers for the encoder and decoder. We denote  $vq\_loss$  and  $mse\_loss$  as  $vq$  and  $mse$ , respectively. The former can be

used to characterize the distribution of latent vectors in the codebook. The smaller  $vq$  indicates a more uniform distribution and better representation of latent vectors in the codebook. The latter represents the overall loss of the reconstructed vectors. Besides, we also calculate the sum of the top 100 values of  $mse\_loss$ , denoted as  $mse\_top100$ . It can effectively assist us in observing the outlier situation of the reconstructed vectors. We can observe that as the number of model layers increases, both  $mse\_loss$  and  $vq\_loss$  decrease

Ratio	Method	Avg.bits	WikiText-2↓	C4↓
–	Llama 2-7B	32	5.12	6.63
~ 8×	GPTQ	4.00	5.49	7.20
	SpQR	3.98	5.28	6.87
	QuIP#	4.02	5.29	6.86
	DuQuant	4.00	6.08	7.79
	AQLM	4.04	5.21	6.75
	VPTQ	4.01	5.64	7.13
	QTIP	4.01	<b>5.17</b>	<b>6.69</b>
	PocketLLM	3.98	5.27	6.86
	SoBP*	11.20	8.64	-
	Wanda*	8.00	6.42	-
DuQuant*	4.01	6.28	7.90	
PocketLLM	3.98	<b>5.74</b>	<b>7.49</b>	

Table 3. The perplexity of quantized Llama 2-7B models with weight updation under 8× compression ratio. – represents the second-best result.

significantly. However, as the number of layers increases further,  $vq\_loss$  significantly increases. We believe that with the further growth of the MLP layer, too many nonlinear factors have been added, making it hard to learn the hidden features of the original vectors well.

Method	layer	rate	MMLU (5)	HellaSwag (0)
Llama 2-7B	–	–	45.34	56.69
PocketLLM	q	8.3%	45.84	56.43
	k	8.3%	46.10	56.12
	q,k	16.6%	45.72	56.35
	v	8.3%	45.66	57.23
	o	8.3%	45.32	55.82
	q,k,v,o	33.2%	44.14	55.93
	gate	22.3%	45.71	56.63
	up	22.3%	45.63	55.89
	down	22.3%	45.62	54.55
	gate,up,down	66.8%	43.53	55.74
all	100%	41.50	55.01	

Table 4. The ablation study on Llama 2-7B models that compress different type of layers for 8× compression ratio.

mlp_layers	vq	mse	mse_top100
1	4.67	4.8e-5	1.26
2	0.46	1.1e-5	0.12
3	0.06	8.9e-6	0.08
5	0.20	9.2e-6	0.09

Table 5. The ablation study on Llama 2-7B models that construct MLP layers in encoder and decoder.

In the Table 6, we observe the impact of codebooks with different sizes on the results. As the size increases, the loss will gradually decrease. When the size of the codebook is smaller than 4096, reducing the size will significantly increase the losses. While the size is greater than 32768, the downward trend of losses will slow down.

In the Table 7, we analyze the effectiveness of RLN and codebook initialization. According to the results, we can find that RLN can effectively reduce the  $mse\_loss$  and improve the reconstruction effect. Codebook initialization can signif-

codebook_size	vq	mse	mse_top100
256	5.54	8.6e-5	0.36
1024	3.12	4.2e-5	0.22
4096	2.41	2.2e-5	0.16
16384	0.38	1.3e-5	0.10
32768	0.06	8.9e-6	0.08

Table 6. The ablation study on Llama 2-7B models with different codebook size.

icantly reduce the  $vq\_loss$  and relieve the outlier situation.

RLN	initial	vq	mse	mse_top100
		8.51	4.5e-5	0.89
	✓	2.36	2.2e-5	0.36
✓		1.23	9.2e-6	0.25
✓	✓	0.06	8.9e-6	0.08

Table 7. The ablation study on the proposed RLN module and codebook initialization.

## Visualization Results

As shown in Figure 3, we visualize the reconstruction results of  $q$ ,  $up$ ,  $down$  weight vectors for different compression ratios. We visualize  $1 \times 4$  weight vectors with the quantity of 16 for  $q$  layers. We visualize  $1 \times 8$  weight vectors with the quantity of 8 for  $up$  and  $down$  layers. As the compression ratio increases, the difference between the original vector and the recovered vector gradually becomes apparent in the details but still remains overall consistent.




Figure 3: The reconstructed results of  $q$ ,  $up$ ,  $down$  weight vectors correspond to the compression ratios of 8× (4 bits), 16× (2 bits), and 20× (1.6 bits), respectively.

## Conclusion

In this paper, we propose PocketLLM, a method for compressing the weights of LLMs at extreme compression ratios. By employing encoder-decoder meta-networks, PocketLLM effectively projects the original weights into a discrete latent space, incorporating nonlinear factors for better representation. It then learns a compact codebook with strong representational power, mapping the codebook’s representative vectors back to the original space. This approach significantly improves upon existing low-bit compression techniques. At 8× compression, PocketLLM without fine-tuning, outperforms current fine-tuning methods and achieves performance comparable to the dense model. At 10× compression, PocketLLM with fine-tuning maintains the dense model’s accuracy, and even at 20× compression, simple fine-tuning restores accuracy.

## References

- Ashkboos, S.; Croci, M. L.; do Nascimento, M. G.; Hoefler, T.; and Hensman, J. 2024. SliceGPT: Compress Large Language Models by Deleting Rows and Columns. *arXiv:2401.15024*.
- Baalen, M. V.; Kuzmin, A.; Nagel, M.; Couperus, P.; Bolshakov, A.; Bastoul, C.; Mahurin, E.; Blankevoort, T.; and Whatmough, P. 2024. GPTVQ: The Blessing of Dimensionality for LLM Quantization. In *Workshop on Efficient Systems for Foundation Models II @ ICML2024*.
- Bengio, Y.; Léonard, N.; and Courville, A. 2013. Estimating or propagating gradients through stochastic neurons for conditional computation. *arXiv preprint arXiv:1308.3432*.
- Boratko, M.; Padigela, H.; Mikkilineni, D.; Yuvraj, P.; Das, R.; McCallum, A.; Chang, M.; Fokoue-Nkoutche, A.; Kapanipathi, P.; Mattei, N.; et al. 2018. A systematic classification of knowledge, reasoning, and context within the ARC dataset. *arXiv preprint arXiv:1806.00358*.
- Chee, J.; Cai, Y.; Kuleshov, V.; and Sa, C. D. 2023. QuIP: 2-Bit Quantization of Large Language Models With Guarantees. *arXiv:2307.13304*.
- Deng, J.; Li, S.; Wang, C.; Gu, H.; Shen, H.; and Huang, K. 2024. LLM-Codebook for Extreme Compression of Large Language Models.
- Dettmers, T.; Svirschevski, R.; Egiazarian, V.; Kuznedelev, D.; Frantar, E.; Ashkboos, S.; Borzunov, A.; Hoefler, T.; and Alistarh, D. 2023. SpQR: A Sparse-Quantized Representation for Near-Lossless LLM Weight Compression. *arXiv preprint arXiv:2306.03078*.
- Egiazarian, V.; Panferov, A.; Kuznedelev, D.; Frantar, E.; Babenko, A.; and Alistarh, D. 2024. Extreme Compression of Large Language Models via Additive Quantization. *arXiv:2401.06118*.
- Frantar, E.; and Alistarh, D. 2023. SparseGPT: Massive Language Models Can Be Accurately Pruned in One-Shot. *arXiv preprint arXiv:2301.00774*.
- Frantar, E.; Ashkboos, S.; Hoefler, T.; and Alistarh, D. 2022. GPTQ: Accurate Post-Training Quantization for Generative Pre-trained Transformers. *arXiv preprint arXiv:2210.17323*.
- Gunter, T.; Wang, Z.; Wang, C.; Pang, R.; Narayanan, A.; Zhang, A.; Zhang, B.; Chen, C.; Chiu, C.-C.; Qiu, D.; et al. 2024. Apple intelligence foundation language models. *arXiv preprint arXiv:2407.21075*.
- Guo, J.; Wu, J.; Wang, Z.; Liu, J.; Yang, G.; Ding, Y.; Gong, R.; Qin, H.; and Liu, X. 2024. Compressing large language models by joint sparsification and quantization. In *Forty-first International Conference on Machine Learning*.
- Hendrycks, D.; Burns, C.; Basart, S.; Zou, A.; Mazeika, M.; Song, D.; and Steinhardt, J. 2020. Measuring Massive Multitask Language Understanding. *CoRR*, abs/2009.03300.
- Hendrycks, D.; and Gimpel, K. 2016. Gaussian Error Linear Units (GELUs). *arXiv preprint arXiv:1606.08415*.
- Hu, E. J.; Shen, Y.; Wallis, P.; Allen-Zhu, Z.; Li, Y.; Wang, S.; Wang, L.; and Chen, W. 2022. LoRA: Low-Rank Adaptation of Large Language Models. In *International Conference on Learning Representations*.
- Huang, X.; Huang, Y.-L.; and Wen, Z. 2025. SoLA: Leveraging Soft Activation Sparsity and Low-Rank Decomposition for Large Language Model Compression. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 17494–17502.
- Ji, Y.; Xiang, Y.; Li, J.; Xia, Q.; Ye, Z.; Duan, X.; Wang, Z.; Chen, K.; and Zhang, M. 2024. Adaptive Feature-based Low-Rank Compression of Large Language Models via Bayesian Optimization. In Al-Onaizan, Y.; Bansal, M.; and Chen, Y.-N., eds., *Findings of the Association for Computational Linguistics: EMNLP 2024*, 4152–4168. Miami, Florida, USA: Association for Computational Linguistics.
- Kaplan, J.; McCandlish, S.; Henighan, T.; Brown, T. B.; Chess, B.; Child, R.; Gray, S.; Radford, A.; Wu, J.; and Amodei, D. 2020. Scaling laws for neural language models. *arXiv preprint arXiv:2001.08361*.
- Kim, S.; Hooper, C.; Gholami, A.; Dong, Z.; Li, X.; Shen, S.; Mahoney, M. W.; and Keutzer, K. 2023. SqueezeLLM: Dense-and-Sparse Quantization. *arXiv preprint arXiv:2306.07629*.
- Kurtic, E.; Kuznedelev, D.; Frantar, E.; Goin, M.; and Alistarh, D. 2023. Sparse Fine-tuning for Inference Acceleration of Large Language Models. *arXiv:2310.06927*.
- Li, J.; Xu, Y.; Huang, H.; Yin, X.; Li, D.; Ngai, E. C.; and Barsoum, E. 2025. Gumiho: A Hybrid Architecture to Prioritize Early Tokens in Speculative Decoding. *arXiv preprint arXiv:2503.10135*.
- Liao, H.; Xu, Y.; He, S.; Li, G.; Yin, X.; Li, D.; Barsoum, E.; Zhao, J.; and Liu, K. 2025. SparK: Query-Aware Unstructured Sparsity with Recoverable KV Cache Channel Pruning. *arXiv preprint arXiv:2508.15212*.
- Lin, H.; Xu, H.; Wu, Y.; Cui, J.; Zhang, Y.; Mou, L.; Song, L.; Sun, Z.; and Wei, Y. 2024. DuQuant: Distributing Outliers via Dual Transformation Makes Stronger Quantized LLMs. *arXiv preprint arXiv:2406.01721*.
- Liu, Y.; Wen, J.; Wang, Y.; Ye, S.; Zhang, L. L.; Cao, T.; Li, C.; and Yang, M. 2024. VPTQ: Extreme Low-bit Vector Post-Training Quantization for Large Language Models. In Al-Onaizan, Y.; Bansal, M.; and Chen, Y.-N., eds., *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, 8181–8196. Miami, Florida, USA: Association for Computational Linguistics.
- Ma, X.; Fang, G.; and Wang, X. 2023. LLM-Pruner: On the Structural Pruning of Large Language Models. In *Advances in Neural Information Processing Systems*.
- Merity, S.; Xiong, C.; Bradbury, J.; and Socher, R. 2016. Pointer sentinel mixture models. *arXiv preprint arXiv:1609.07843*.
- OpenAI. 2023. ChatGPT (Mar 14 version). <https://chat.openai.com/chat>.
- Raffel, C.; Shazeer, N.; Roberts, A.; Lee, K.; Narang, S.; Matena, M.; Zhou, Y.; Li, W.; and Liu, P. 2020. Exploring the Limits of Transfer Learning with a Unified Text-to-Text Transformer. *Journal of Machine Learning Research*, 21(140): 1–67.

- Sakaguchi, K.; Bras, R. L.; Bhagavatula, C.; and Choi, Y. 2021. WinoGrande: an adversarial winograd schema challenge at scale. *Commun. ACM*, 64(9): 99–106.
- Shao, W.; Chen, M.; Zhang, Z.; Xu, P.; Zhao, L.; Li, Z.; Zhang, K. Z.; Gao, P.; Qiao, Y.; and Luo, P. 2023. OmniQuant: Omnidirectionally Calibrated Quantization for Large Language Models. *arXiv preprint arXiv:2308.13137*.
- Sun, M.; Liu, Z.; Bair, A.; and Kolter, J. Z. 2023. A Simple and Effective Pruning Approach for Large Language Models. *arXiv preprint arXiv:2306.11695*.
- Taori, R.; Gulrajani, I.; Zhang, T.; Dubois, Y.; Li, X.; Guestrin, C.; Liang, P.; and Hashimoto, T. B. 2023. Stanford Alpaca: An Instruction-following LLaMA model. [https://github.com/tatsu-lab/stanford\\_alpaca](https://github.com/tatsu-lab/stanford_alpaca).
- Tata, S.; and Patel, J. M. 2003. PiQA: An algebra for querying protein data sets. In *International Conference on Scientific and Statistical Database Management*.
- Touvron, H.; Lavril, T.; Izacard, G.; Martinet, X.; Lachaux, M.-A.; Lacroix, T.; Rozière, B.; Goyal, N.; Hambro, E.; Azhar, F.; et al. 2023a. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*.
- Touvron, H.; Martin, L.; Stone, K.; Albert, P.; Almahairi, A.; Babaei, Y.; Bashlykov, N.; Batra, S.; Bhargava, P.; Bhosale, S.; Bikel, D.; Blecher, L.; Ferrer, C. C.; Chen, M.; Cucurull, G.; Esiobu, D.; Fernandes, J.; Fu, J.; Fu, W.; Fuller, B.; Gao, C.; Goswami, V.; Goyal, N.; Hartshorn, A.; Hosseini, S.; Hou, R.; Inan, H.; Kardas, M.; Kerkez, V.; Khabsa, M.; Kloumann, I.; Korenev, A.; Koura, P. S.; Lachaux, M.-A.; Lavril, T.; Lee, J.; Liskovich, D.; Lu, Y.; Mao, Y.; Martinet, X.; Mihaylov, T.; Mishra, P.; Molybog, I.; Nie, Y.; Poulton, A.; Reizenstein, J.; Rungta, R.; Saladi, K.; Schelten, A.; Silva, R.; Smith, E. M.; Subramanian, R.; Tan, X. E.; Tang, B.; Taylor, R.; Williams, A.; Kuan, J. X.; Xu, P.; Yan, Z.; Zarov, I.; Zhang, Y.; Fan, A.; Kambadur, M.; Narang, S.; Rodriguez, A.; Stojnic, R.; Edunov, S.; and Scialom, T. 2023b. Llama 2: Open Foundation and Fine-Tuned Chat Models. *arXiv:2307.09288*.
- Tseng, A.; Chee, J.; Sun, Q.; Kuleshov, V.; and Sa, C. D. 2024a. QuIP#: Even Better LLM Quantization with Hadamard Incoherence and Lattice Codebooks. *arXiv:2402.04396*.
- Tseng, A.; Sun, Q.; Hou, D.; and De, C. 2024b. QTIP: Quantization with Trellises and Incoherence Processing. In Globerson, A.; Mackey, L.; Belgrave, D.; Fan, A.; Paquet, U.; Tomczak, J.; and Zhang, C., eds., *Advances in Neural Information Processing Systems*, volume 37, 59597–59620. Curran Associates, Inc.
- Wang, J.; Zeng, Y.; Guo, J.; Ma, Y.; Liu, A.; and Liu, X. 2025. SLMQuant: Benchmarking Small Language Model Quantization for Practical Deployment. In *3rd International Workshop on Rich Media With Generative AI*.
- Weber, M.; Fu, D. Y.; Anthony, Q.; Oren, Y.; Adams, S.; Alexandrov, A.; Lyu, X.; Nguyen, H.; Yao, X.; Adams, V.; Athiwaratkun, B.; Chalamala, R.; Chen, K.; Ryabinin, M.; Dao, T.; Liang, P.; Ré, C.; Rish, I.; and Zhang, C. 2024. RedPajama: an Open Dataset for Training Large Language Models. *NeurIPS Datasets and Benchmarks Track*.
- Wei, J.; Lu, Q.; Jiang, N.; Li, S.; Xiang, J.; Chen, J.; and Liu, Y. 2024. Structured Optimal Brain Pruning for Large Language Models. In Al-Onaizan, Y.; Bansal, M.; and Chen, Y.-N., eds., *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, 13991–14007. Miami, Florida, USA: Association for Computational Linguistics.
- Xu, Y.; Xie, L.; Gu, X.; Chen, X.; Chang, H.; Zhang, H.; Chen, Z.; Zhang, X.; and Tian, Q. 2023. QA-LoRA: Quantization-Aware Low-Rank Adaptation of Large Language Models. *arXiv:2309.14717*.
- Xu, Z.; Zhang, Y.; Xie, E.; Zhao, Z.; Guo, Y.; Wong, K.-Y. K.; Li, Z.; and Zhao, H. 2024. Drivegpt4: Interpretable end-to-end autonomous driving via large language model. *IEEE Robotics and Automation Letters*.
- Yang, A.; Li, A.; Yang, B.; Zhang, B.; Hui, B.; Zheng, B.; Yu, B.; Gao, C.; Huang, C.; Lv, C.; Zheng, C.; Liu, D.; Zhou, F.; Huang, F.; Hu, F.; Ge, H.; Wei, H.; Lin, H.; Tang, J.; Yang, J.; Tu, J.; Zhang, J.; Yang, J.; Yang, J.; Zhou, J.; Zhou, J.; Lin, J.; Dang, K.; Bao, K.; Yang, K.; Yu, L.; Deng, L.; Li, M.; Xue, M.; Li, M.; Zhang, P.; Wang, P.; Zhu, Q.; Men, R.; Gao, R.; Liu, S.; Luo, S.; Li, T.; Tang, T.; Yin, W.; Ren, X.; Wang, X.; Zhang, X.; Ren, X.; Fan, Y.; Su, Y.; Zhang, Y.; Zhang, Y.; Wan, Y.; Liu, Y.; Wang, Z.; Cui, Z.; Zhang, Z.; Zhou, Z.; and Qiu, Z. 2025. Qwen3 Technical Report. *arXiv preprint arXiv:2505.09388*.
- Yang, H.; Liu, X.-Y.; and Wang, C. D. 2023. Fingpt: Open-source financial large language models. *arXiv preprint arXiv:2306.06031*.
- Yao, Z.; Aminabadi, R. Y.; Zhang, M.; Wu, X.; Li, C.; and He, Y. 2022. ZeroQuant: Efficient and Affordable Post-Training Quantization for Large-Scale Transformers. *arXiv preprint arXiv:2206.01861*.
- Zellers, R.; Holtzman, A.; Bisk, Y.; Farhadi, A.; and Choi, Y. 2019. HellaSwag: Can a Machine Really Finish Your Sentence? In Korhonen, A.; Traum, D. R.; and Màrquez, L., eds., *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, 4791–4800. Association for Computational Linguistics.
- Zeng, A.; Liu, X.; Du, Z.; Wang, Z.; Lai, H.; Ding, M.; Yang, Z.; Xu, Y.; Zheng, W.; Xia, X.; Tam, W. L.; Ma, Z.; Xue, Y.; Zhai, J.; Chen, W.; Liu, Z.; Zhang, P.; Dong, Y.; and Tang, J. 2023. GLM-130B: An Open Bilingual Pre-trained Model. In *The Eleventh International Conference on Learning Representations (ICLR)*.