

# Agile-Quant: Activation-Guided Quantization for Faster Inference of LLMs on the Edge

Xuan Shen<sup>\*1</sup>, Peiyan Dong<sup>\*1</sup>, Lei Lu<sup>1</sup>, Zhenglun Kong<sup>1</sup>,  
Zhengang Li<sup>1</sup>, Ming Lin<sup>†2</sup>, Chao Wu<sup>1</sup>, Yanzhi Wang<sup>1</sup>

<sup>1</sup>Northeastern University

<sup>2</sup>Oracle

{shen.xu, dong.pe, lu.lei1, kong.zhe, li.zhen, cha.wu, yanz.wang}@northeastern.edu  
linming04@gmail

## Abstract

Large Language Models (LLMs) stand out for their impressive performance in intricate language modeling tasks. However, their demanding computational and memory needs pose obstacles for broad use on edge devices. Quantization is then introduced to boost LLMs' on-device efficiency. Recent works show that 8-bit or lower weight quantization is feasible with minimal impact on end-to-end task performance, while the activation is still not quantized. On the other hand, mainstream commodity edge devices still struggle to execute these sub-8-bit quantized networks effectively. In this paper, we propose Agile-Quant, an Activation-Guided quantization framework for faster Inference of popular Large Language Models (LLMs) on the Edge. Considering the hardware profiling and activation analysis, we first introduce a basic activation quantization strategy to balance the trade-off of task performance and real inference speed. Then we leverage the activation-aware token pruning technique to reduce the outliers and the adverse impact on attentivity. Ultimately, we utilize the SIMD-based 4-bit multiplier and our efficient TRIP matrix multiplication to implement the end-to-end accelerator for LLMs on multiple edge devices. We apply our framework on different scales of LLMs including LLaMA, OPT, and BLOOM with 4-bit or 8-bit for the activation and 4-bit for the weight quantization. Experiments show that Agile-Quant achieves simultaneous quantization of model weights and activations while maintaining task performance comparable to existing weight-only quantization methods. Moreover, in the 8- and 4-bit scenario, Agile-Quant achieves an on-device speedup of up to 2.55x compared to its FP16 counterparts across multiple edge devices, marking a pioneering advancement in this domain.

## Introduction

Large Language Models (LLMs) (Touvron et al. 2023; Zhang et al. 2022; Brown et al. 2020a; Radford et al. 2019; Brown et al. 2020b) based on the Transformer (Vaswani et al. 2017) family have breakthrough performance in Natural Language Processing (NLP) research area.

**Application Scenarios.** In real-world decision scenarios, incorporating LLMs inference as a crucial element often

<sup>\*</sup>These authors contributed equally.

<sup>†</sup>Work done before joining Oracle.

Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

necessitates stringent latency requirements. However, one drawback of LLMs is their computational and storage cost, which ranks among the highest for known models. Consider GPT3-175B as an example. When stored in a compact float16 format, its parameters require 326GB (in multiples of 1024) of memory. This surpasses the capacity of even the most powerful individual GPUs, not to mention the challenges of running it on hardware-limited edge devices with acceptable latency. Quantization, in particular, offers a promising approach to substantially improve the inference throughput and energy efficiency of LLMs on edge devices. This improvement is achieved by harnessing the highly effective 8-bit fixed-point (INT8) operations supported by the SIMD units that are commonly found in edge platforms, such as CPUs and Raspberry Pis.

**Current Limitations.** Before fully realizing the on-device benefits of model quantization on LLMs, it's crucial to address two pressing issues that demand careful attention. ❶ Existing works (Frantar et al. 2022; Lin et al. 2023; Xiao et al. 2022) primarily concentrate on weight-only (4-bit) quantization while leaving activations in the floating-point (FP16) domain. This approach limits the efficient speed-up of model inference on common edge devices, which typically only support 16x16 and 8x8 integer multipliers. Specifically, activation quantization often has a detrimental effect on task performance, especially when the model size becomes large, due to the emergence of pronounced outliers in activations. Experiments done by work (Dettmers et al. 2022) indicate that directly setting these outliers to zero can result in a substantial 45% degradation in task performance. Additionally, given the large model size of LLMs, limited academic computing power makes it challenging to afford the associated training costs. Consequently, Post-Training Quantization (PTQ) has become a prevalent approach, but it falls short of minimizing the quantization error caused by these outliers. In summary, quantizing the activations of LLMs while handling outliers inside activations is a crucial yet challenging issue. ❷ Mainstream edge processors, such as CPUs and Raspberry Pis, leverage SIMD units to execute multiple operations in parallel efficiently. SIMD instructions are adept at exploiting byte-level data (8-bit integers) parallelism and are well-supported in common ISAs (Instruction Set Architectures) and DNN processing frameworks. Examples include GEMMLOWP (Jacob and Warden 2017)

in TensorFlow Lite and QNNPACK (Dukhan, Wu, and Lu 2018) in PyTorch. Their low-precision kernels merely zero-extend the sub-byte operands to align them with byte boundaries, treating them as 8-bit or 16-bit operands.

In this paper, we address the above on-device quantization issues while enjoying the powerful performance provided by LLMs. We propose Agile-Quant, an activation-guided quantization framework for faster inference of LLMs on the edge. Specifically, we begin with a fundamental activation quantization strategy based on hardware latency profiling and activation analysis of LLMs, aiming to strike a balance between task performance and on-device inference speed. We subsequently utilize the activation-aware pruning method to optimize quantization. This is crucial because quantized tokens often exhibit numerous outliers, causing their attention to shift from the first position to nearby local positions. By pruning tokens, we effectively eliminate some outliers, as they typically concentrate within the same or adjacent channels of different tokens. Also, the removal of inattentive tokens can reduce the interaction distance between important tokens. Finally, we design the edge-oriented optimization for the hardware implementation of Agile-Quant. It consists primarily of two components: a SIMD-based 4-bit multiplier to facilitate efficient 4x4 INT4 multiplication, and our efficient Two-Refine Improved by Pruning (TRIP) matrix multiplication designed to mitigate the adverse impact of outliers.

The popular LLMs models such as LLaMA (Touvron et al. 2023), OPT (Zhang et al. 2022), and BLOOM (Scao et al. 2022) are adopted to verify the effectiveness of our framework and the efficiency of our method on multiple edge devices. Agile-Quant can maintain state-of-the-art task performance comparable with weight-only works while achieving practical on-device speedup up to 2.55x.

The contributions of this work are summarized as follows:

- We design the activation-guided and edge-oriented quantization strategy for the balance of latency decreasing and task performance.
- We design an activation-aware token pruning method to minimize the negative impact on task performance caused by the outliers and the local attentivity.
- We propose the SIMD-based 4-bit multiplier and an efficient TRIP matrix multiplication for effective hardware implementation.
- We achieve state-of-the-art task performance on several popular datasets with practical on-device speedup.

## Background and Related Works

In this section, we first focus on the background of post-training quantization for LLMs. Then we discuss the low-bit computation on general edge devices.

### Post-Training Quantization for LLMs

Post-Training Quantization (PTQ) techniques are widely used for one-shot compressing models, particularly for Large Language Models (LLMs), given the high cost of re-training. These PTQ methods employ accurate solvers to address compression challenges on a per-layer or per-group basis, relying on a limited set of calibration data. Notably,

recent advances in PTQ, like GPTQ (Frantar et al. 2022), AWQ (Lin et al. 2023), and SpQR (Dettmers et al. 2023), have introduced well-crafted approaches capable of preserving LLM performance effectively. GPTQ leverages second-order information to correct errors, achieving commendable accuracy within a 3-4 bit range. AWQ proposes safeguarding only 1% of crucial weights to substantially diminish quantization errors. SpQR’s focus is on reducing quantization to 3-4 bits per parameter for smaller models. Moreover, they put forth a novel technique enabling nearly lossless compression of LLMs. Nonetheless, these works fall short of achieving practical inference acceleration on edge devices, as the activation part persists in a floating-point format, rendering the integer multiplier of the edge devices ineffective.

### Low-Bit Computation on Hardware Devices

Low-precision linear algebra kernels aim to maximize computing throughput on low-precision operands. This is achieved by extending existing wider bit-width linear algebra kernels. The use of lower-precision operands brings about two performance enhancements: increased cache capacity and the ability to leverage lower-precision SIMD instructions for processing multiple elements simultaneously. Pioneering examples of these low-precision linear algebra kernels, e.g., Google’s GEMMLOWP (Jacob and Warden 2017) and Facebook’s QNNPACK (Dukhan, Wu, and Lu 2018), excel at enhancing the efficiency of DNN inference when employing 8-bit quantization. However, pushing for more aggressive sub-byte quantization yields no added performance benefits due to the fact that mainstream CPUs solely support SIMD operations with a precision of 8 bits or wider. In specific, low-precision kernels essentially expand sub-byte operands to 8 bits and process them accordingly. Furthermore, the concept of Bit-serial computation emerges as a promising solution for data-parallel computation with sub-byte values. This approach involves sequentially processing each bit of two operands during multiplication, while simultaneously managing multiple operand pairs in parallel. Nonetheless, its practical implementation necessitates the *popcount* operation, which inherently limits runtime throughput. As a result, this method only presents significant advantages in ultra-low-bit scenarios (1 or 2 bits).

### Activation Analysis of LLMs

In this section, we analyze the attentivity of tokens in LLMs and the influence of token pruning on activation quantization. Besides, we deliver the latency profiling to analyze potential quantization strategy.

### Token Importance in LLMs

In natural language processing, numerous non-essential words often exist within sentences, contributing little to the overall comprehension. This implies that we can efficiently process these words using fewer resources, potentially even excluding them, in order to mitigate complexity. As words are embedded into tokens in language models, we explore the attention mechanism to analyze the importance of each token. The previous works (Kong et al. 2022; Dong

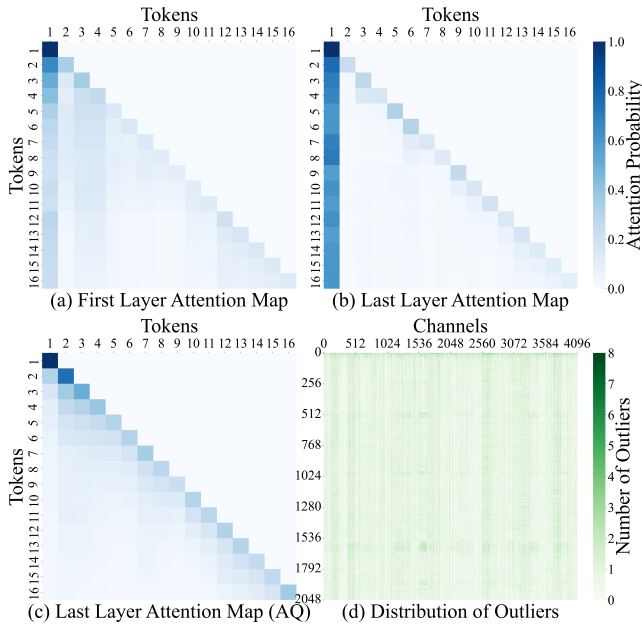


Figure 1: The (a), (b), and (c) shows attention maps with 16 tokens in the first and last layer of the model. The activation is not quantized in (a) and (b), while it is quantized in (c). The (d) shows the distribution of outliers in one activation with 2048 tokens. The visualization is based on the LLaMA-7B model with the Wikitext-2 dataset.

et al. 2023) focus on the attention map in the transformer architectures. The attention probabilities are then accumulated across multiple rounds of attention as token importance scores. However, the causal attention masks used in LLMs ensure that, during the self-attention mechanism, each token can only interact with previous tokens instead of the following ones. Thus, this causal mechanism makes the accumulated probabilities not appropriate to the evaluation of token importance because of its unfair for the accumulated probabilities of former tokens.

In LLMs, a distinct start token is placed at the beginning of the input sequence. The start token has a role in initializing the hidden layers and defining token positions within the sequence. These aspects are vital for producing text that is both coherent and contextually meaningful. To explore the relationship between the first start token and other tokens, we visualize the attention map at the first and last layer of the LLaMA-7B model with 16 tokens on the Wikitext-2 dataset in Figure 1 (a) and (b). According to the attention map, several tokens in the first layer demonstrate a shared triangular pattern, indicating that tokens tend to the adjacent positions, especially the previous position. While in the last layer, nearly all tokens share a vertical-stripe pattern, indicating that tokens all related with the first token. Then we explore the attention maps in the middle layers, showing that these maps are similar to the one in the last layer. Thus, it guides us to build the connection between the token importance and token attentivity to the start token.




Figure 2: Mobile Device profiling of one LLaMA block.

### Influence of Activation Quantization

We analyze the distribution of outliers and visualize outlier in different channels in Figure 1 (d). We notice that the outliers are distributed in adjacent or even the same channels, since several straight lines with deep colors indicates that the channel index of the outliers unchange. Also, the attention map, which is generated by the query  $Q$  and key matrix  $K$ , can be influenced by the activation quantization as it is input-dependent. We visualize the quantized attention map at the last layer in Figure 1 (c). The attention map shows a triangular pattern and the quantized tokens attend to the adjacent positions rather than the start token, demonstrating that the attention range becomes locality and the attentivity turns much weaker. This change implies a deterioration in the globality of representative features. From another perspective, the information reduction of the original attention map caused by quantization error will impact the final task performance adversely.

### Latency Profiling on Hardware Devices

To gain a deeper insight into the runtime distribution of LLMs, we conducted profiling on a widely used model, LLaMA, utilizing the on-board Snapdragon 870 CPU, as shown in Figure 2. This profiling includes FP16, INT8, and INT4 precisions. Since nonlinear operators (LayerNorm/Softmax/SwiGLU) contribute a relatively smaller portion of latency, i.e.,  $< 8\%$  for FP16,  $< 12\%$  for INT8,  $< 16\%$  for INT4, we have implemented them using FP16 arithmetic units to ensure task performance is maintained. We focus on the primary computation workload, matrix multiplication, performed in various low-bit precision settings. Our observation reveals that FC1 and FC2 account for 54% of the runtime latency in FP16, and 49.5% in INT8. This finding indicates the need to prioritize the quantization of these components to a lower-bit (4-bit) representation. Following that, our order of priority will be as follows: Linear Transformation  $>$  Linear Projection  $>$  AttnV  $>$  QK. In essence, focusing on low-bit quantization of both weights and activations for LLMs while ensuring task performance is crucial.

### Methodology

We explain the activation quantization pipeline here and propose the activation-guided framework for the optimization of quantization. Also, we explain our hardware implementation of the 4-bit multiplier.

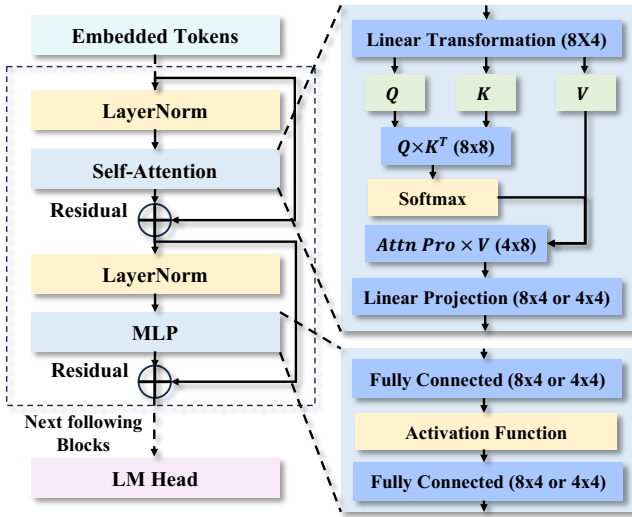


Figure 3: Activation Quantization Pipeline.

### Preliminary

We here explain the quantizers we use for activation quantization in this work. We assume the bit-width used in quantization is  $b$ , and then the quantizer can be defined as a function  $Q(X|b)$  which can map the floating points in vector  $X \in \mathbb{R}^{m \times n}$  to the closest quantization in  $q$ :

$$q = \begin{cases} \{-2^{b-1}, \dots, 2^{b-1} - 1\}, & \text{Signed} \\ \{0, 1, \dots, 2^b - 1\}, & \text{Unsigned} \end{cases} \quad (1)$$

There are various kinds of quantizers  $Q(X|b)$ , and the uniform quantizer (Jacob et al. 2018) and the log2 quantizer (Cai, Takemoto, and Nakajo 2018) are widely used. In our work, we mainly use these two quantizers for activation quantization.

**Uniform Quantization** has been supported by most hardware devices.

$$Q(X|b) = \text{CLIP}(\lfloor \frac{X}{s} \rfloor + zp, 0, 2^b - 1) \quad (2)$$

The  $s$  and  $zp$  denote the scale and zero-point separately.

**Log2 Quantization** imports the exponential operation into the linear quantization process.

$$Q(X|b) = \text{Sign}(X) \cdot \text{CLIP}(\lfloor -\log_2 \frac{X}{\max(|X|)} \rfloor, 0, 2^{b-1} - 1) \quad (3)$$

### Activation Quantization Pipeline

We present our activation quantization pipeline in Figure 3. While the embedding process, output module, and the yellow-highlighted nonlinear operations contribute relatively small proportions during model inference, we preserve their computation without alteration. Our primary focus is optimizing the matrix multiplication operations, constituting the largest share of the inference latency.

Within our pipeline, we target the acceleration of computations occurring in the self-attention and MLP modules, as indicated by the blue shading. Specifically, we perform activation quantization predominantly using 8-bit integers.




Figure 4: Activation Quantization With Token Pruning.

However, we observe that specific activations following the self-attention mechanism can be quantized using 4-bit integers, resulting in further acceleration while upholding task performance standards. Accordingly, we propose our innovative 4-bit multiplier to effectively support the INT4 matrix multiplication.

### Activation-Guided Optimization

Based on the analysis of outliers and attention range in the previous section, it is intuitive for us to import the token pruning here for optimization. We visualize the token pruning process in Figure 4. Token pruning can reduce the outliers, which can decrease the quantization error caused by them. Token pruning can also reduce the distance between attentive tokens to help the model capture more features.

We first introduce the activation-aware token pruning improved activation quantization method we use in this work. Inspired by the work (Lin et al. 2022), we propose the Two-Refine Improved by Pruning (TRIP) method here to address the difficulties in activation quantization.

For the one activation in transformer-based models, we assume it as  $X \in \mathbb{R}^{m \times d}$ , and the  $m, d$  denote the number of tokens, and dimension separately. We assume the token pruning function, which prunes tokens in cascade according to the token importance, as  $F^P(\cdot)$ .

$$X^P = F^P(X) \in \mathbb{R}^{n \times d}, n < m \quad (4)$$

Then, the TRIP factor  $\alpha = \{\alpha_1, \alpha_2, \dots, \alpha_d\}$  are applied to the different channels whose number of outliers is reduced by token pruning. The factor  $\alpha$  can provide different channels with different factors, which can regularize the quantization parameters and is utilized as:

$$X_Q^P = Q(X^P|b, \alpha) = \text{CLIP}(\lfloor \frac{X^P}{2^{\alpha_s}} \rfloor + zp, 0, 2^b - 1) \quad (5)$$

$$s = \frac{\max(X^P) - \min(X^P)}{(2^b - 1) \cdot 2^\alpha} \quad (6)$$

$$zp = \text{CLIP}(\lfloor -\frac{\min(X^P)}{2^\alpha \cdot s} \rfloor, 0, 2^b - 1) \quad (7)$$

For channel  $c \in [1, d]$  with the biggest refinement  $K$ , the factor  $\alpha_c$  is as follow:

$$\alpha_c = \arg \min_{\alpha_c \in \{0, 1, \dots, K\}} \left\| X_c^P - \lfloor \frac{X_c^P}{2^{\alpha_c \cdot s}} \rfloor \cdot 2^{\alpha_c \cdot s} \right\|_2 \quad (8)$$

The  $K$  is dependent on the channels containing relatively more outliers. Outliers can be reduced by token pruning.




Figure 5: The paradigm of INT4 multiplier.

For token pruning, we evaluate the token importance based on the attention map, and we only adopt their attentivity to the first start token for the importance measuring based on our analysis. We prune tokens in cascade at the different depths of the model progressively and dynamically. Inspired by the previous token pruning works (Kim et al. 2022; Liang et al. 2022), we design our pruning strategy by splitting the model into several stages which have different redundancy (Shen et al. 2023). Meanwhile, pruning tokens from some shallow layer or even from the first layer would influence the task performance of the model. It is because, at the first few layers, the attentivity of tokens is still locality and the larger number of tokens can help the model capture more beneficial features. Therefore, we mainly adopt the strategy that progressively pruning the tokens starting from the layer whose attention map shows that the tokens have enough ability to capture features. Besides, we regulate the pruning ratio in each layer to balance the trade-off between loss of information and the improvement to the quantization.

### Edge-Oriented Optimization

This section mainly proposes the hardware implementation of our Agile-Quant framework on edge devices. We first design the SIMD-based 4-bit multiplier to support the INT4 multiplication and then introduce the efficient support of 2-refined matrix multiplication.

**INT4 Multiplier.** We designed a specialized 4-bit multiplier based on SIMD architecture (Figure 5), aimed at supporting practical INT4 computation. Here’s the workflow: Dot-multiplication: In the SIMD kernel, we combine two adjacent weight values,  $W_{i,j}$  and  $W_{i+1,j}$ , and multiply them with their shared activation value. The result is an INT16 data type. We allocate the first 8 bits for the multiplication with  $W_{i,j}$  and the remaining 8 bits for  $W_{i+1,j}$ . This approach follows the SIMD memory mechanism. Addition: By utilizing the Bitshift operator, we expand the 16-bit output from Step 1 to 32 bits. The first 8 bits are set to 0, followed by Output $_{i,j}$  in the next 8 bits, 0s in the third 8 bits, and Output $_{i+1,j}$  in the final 8 bits. We then perform a row-by-row summation. This process can handle up to  $2^8$  additions without overflow, sufficient for multi-head attention (head-dimension = 32/64). Each addition has a 32-bit memory footprint. Finally, we split the output into two INT16 values and quantize them back to INT4 at the value-level, allowing us to integrate them into the GeMM kernel.

**Efficient TRIP Matrix Multiplication.** Unlike channel-wise quantization, we perform layer-wise quantization on the activation matrix of outliers. All the channels share the

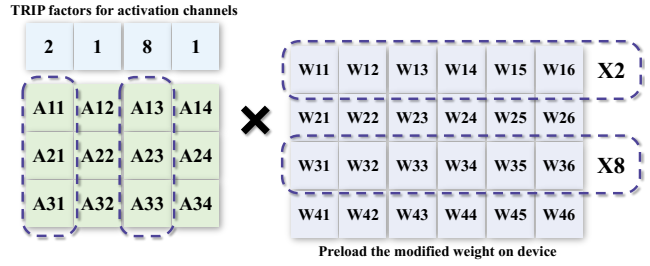


Figure 6: Hardware Implementation of Efficient TRIP Matrix Multiplication.

same quantization parameters, i.e., scaling factors and zero-point. However, the predicted TRIP factors will adapt to the outlier channels’ scaling factors. In the practical implementation, those TRIP factors will be equivalently mathematically transformed over the corresponding weights, as shown in Figure 6.

Note that common inference engines only support layer-wise quantization on activation and per-channel quantization on weights, such as the GeMM and Convolution operator configuration. For example, ArmComputeLibrary (ARM 2023) only supports channel-wise quantization configuration for weight matrix instead of input activation.

## Experiments and Results

We introduce the experiments and the results in this section to verify the effectiveness and efficiency of our method.

### Experiment Setup

**Setup for Activation-guided Quantization.** We implement the activation quantization based on the weight-only quantization work GPTQ (Frantar et al. 2022) which achieves state-of-the-art performance with 4-bit weight-only quantization for LLMs. We mainly use 4-bit and 8-bit integers in our activation quantization. We use the Log2 quantization for softmax activation quantization and use our TRIP quantization for other activations. We implement the different scales of LLaMA, OPT, and BLOOM models in our experiments on the Wikitext-2 dataset (Merity et al. 2016) and C4 (Raffel et al. 2020) dataset.

**Hardware Platform.** We test the actual inference implementation on various edge devices, including the RealmeGT Android Phone with Snapdragon 870 SoC and Raspberry4 B with Quad-core CPU and 8GB RAM. Our inference engine for Arm processors is modified based on ArmComputeLibrary v22.05. The inference latency is reported via the average of 50 iterations for each test.

### Regulation of Token Pruning

We regulate the token pruning ratio to optimize the task performance of LLMs. We apply the prune ratio progressively starting from the shallow layers so that the token pruning can optimize the activation quantization for more deep layers. Assume the model has  $n$  layers with  $L = (l_1, l_2, \dots, l_n)$  and the pruning operation is added to the layers  $L_p = (l_{p1}, l_{p2}, \dots, l_{pm})$ . We set prune ratio  $\beta$  at the last layer  $l_n$




Figure 7: Token Sparsity vs. Perplexity. The visualization is based on LLaMA with 7B, 13B, and 30B scales on the Wikitext-2 dataset.

and compute the progressive ratio  $\gamma$  for the layers  $l_i \in L_p$  as:

$$\gamma = 1 - (1 - \beta)^{\frac{1}{m}} \quad (9)$$

We accumulate the token sparsity  $s$  with the number of pruned tokens during inference as:

$$s = 1 - \sum_{i=1}^n r_i \quad (10)$$

The  $r_i$  denotes the number of remaining tokens in  $l_i$ .

We then adopt the weight and activation both quantized LLaMA models with 7B, 13B, and 30B to search for the optimal prune ratio. We visualize our results in Figure 7. For all three different scales of LLaMA models, we use token pruning to achieve better performance than dense models at the red star points. Also, we find that token pruning can only help the quantization achieve better results when the token sparsity is small, while token pruning makes a negative impact on the task performance when the token sparsity becomes too large. Here we adopt the optimal token sparsity for different scales of the LLaMA model and show the exact results in Table 1. Also, we regulate the token pruning with the same strategy as LLaMA for OPT and BLOOM models, and the best results are shown in Table 2 and Table 3.

## Quantization Results and Analysis

We first show the quantization results of LLaMA in Table 1. According to the results of weight-only quantization works, our method achieves a minor task performance drop, and we achieve better performance than most of the other activation quantization works. The data in the last row denotes the results achieved by token pruning in our method, which verifies that the token pruning can optimize the quantization. Our method achieves better task performance than MoFQ8 (Zhang et al. 2023) and ZeroQuant-FP (Wu, Yao, and He 2023), the Low-Rank Compensation (LoRC) enhancement denoted as  $\S$  would increase the model size and flops, which is also not well-supported on the common inference engine. Then, we show our quantization result of OPT and BLOOM on the Wikitext-2 dataset and C4 dataset in Table 2 and Table 3. Our method achieve better task performance than those activation quantization works and our results are close to the weight-only quantization works. Mean-

Method	WQ # Bits	AQ # Bits	PPL of LLaMA			
			7B	13B	30B	65B
-	FP16	FP16	5.68	5.09	4.10	3.53
RTN	INT4	FP16	6.29	5.53	4.54	3.92
GPTQ	INT4	FP16	5.85	5.2	4.23	3.65
SqLLM	4.05 $\ddagger$	FP16	5.79	5.18	4.22	-
SpQR	3.94 $\ddagger$	FP16	5.87	5.22	4.25	3.90
MoFQ8	FP8 $\dagger$	FP8 $\dagger$	6.49	5.41	5.31	-
ZQFP	INT8	INT8	5.72	5.09	4.10	-
ZQFP	INT4	INT8	6.44	5.32	4.36	-
ZQFP $\S$	INT4	INT8	5.88	5.28	4.34	-
AgileQ	AgileQ-8		6.16	5.57	4.55	4.01
AgileQ*	AgileQ-8		6.09	5.21	4.44	3.92

Table 1: LLaMA Quantization Results on Wikitext-2 dataset. AgileQ-8 denotes the 8-bit is used. SqLLM denotes SqueezeLLM. \* denotes the token pruning optimized results.  $\dagger$  denotes the mix precision with mainly FP8 and INT8.  $\ddagger$  denotes the average bits. ZQFP denotes ZeroQuant-FP.  $\S$  denotes the LoRC.

while, token pruning also helps us achieve better task performance in activation quantization. Especially, for OPT and BLOOM models, our method can even achieve even better task performance than the FP16 models on the C4 dataset.

## Ablation Study

The length of the input sequence makes a big influence on the evaluation process, we try the token pruning with different sequence lengths to further explore the variation of performance. The normal default length of the input sequence used in the evaluation of LLMs is 2048, which is widely used in LLMs-related works. Thus, we regulate the input sequence length of LLaMA-7B with Wikitext-2 dataset to explore the relationship between it and token sparsity. The results are in Figure 8. We can find that task performance becomes worse as the sequence length becomes shorter and the token sparsity becomes larger, and the token pruning only works when the sequence length is long enough (i.e., 2048).




Figure 8: Token Sparsity vs. Input Sequence Length. The visualization is based on LLaMA-7B with Wikitext-2 dataset.

Method	W / A	PPL of OPT on Wikitext-2						PPL of OPT on C4					
	# Bits	125M	1.3B	2.7B	6.7B	13B	30B	125M	1.3B	2.7B	6.7B	13B	30B
-	16	27.65	14.63	12.47	10.86	10.13	9.56	26.56	16.07	14.34	12.71	12.06	11.44
RTN	<b>4/16</b>	37.28	48.17	16.92	12.10	11.32	10.98	33.91	24.51	18.43	14.36	13.36	13.46
GPTQ	<b>4/16</b>	31.12	15.47	12.87	11.39	10.31	9.63	29.22	16.97	15.00	13.18	12.26	11.57
AWQ	<b>4/16</b>	33.96	16.85	14.61	12.44	11.60	10.75	-	-	-	-	-	-
MoFQ8	8 <sup>†</sup> / 8 <sup>†</sup>	-	16.78	14.24	12.41	12.52	10.95	-	-	-	-	-	-
ZQV2	<b>4/16</b>	36.71	19.38	17.92	11.91	10.67	10.10	30.92	17.93	18.32	13.01	12.07	11.33
AgileQ	our-8	31.52	15.90	13.43	11.43	10.42	9.70	28.43	16.72	14.91	12.70	11.77	11.14
AgileQ*	our-8	30.37	14.90	13.19	11.21	10.00	9.45	24.44	15.95	14.20	12.39	11.31	11.20

Table 2: Perplexity of OPT model on Wikitext-2 dataset and C4 dataset. our-8 denotes the 8-bit is used. The bold part denotes integer, otherwise float. \* denotes token pruning optimized results. † denotes mix precision with mainly FP8 and INT8.

Method	W/A	PPL of BLOOM on Wikitext-2					PPL of BLOOM on C4				
	# Bits	560M	1.1B	1.7B	3B	7.1B	560M	1.1B	1.7B	3B	7.1B
-	16/16	22.42	17.69	15.39	13.48	11.37	26.60	22.05	19.49	17.49	15.2
RTN	<b>4 / 16</b>	25.9	22.00	16.97	14.76	12.10	29.89	24.44	21.26	18.76	16.06
GPTQ	<b>4 / 16</b>	24.03	19.05	16.48	14.20	11.73	28.00	23.25	20.55	18.10	15.60
ZQV2	<b>4 / 16</b>	25.31	23.90	16.93	14.65	12.06	27.10	25.99	19.47	17.26	14.83
AgileQ	our-8	24.01	18.82	16.23	14.05	11.73	26.39	21.80	19.18	16.96	14.70
AgileQ*	our-8	23.72	18.33	16.15	13.73	11.36	25.21	19.92	18.56	16.24	14.03

Table 3: Perplexity of BLOOM model on the Wikitext-2 dataset and C4 dataset. our-8 denotes the 8-bit is used. The bold part denotes integer, otherwise float. \* denotes the token pruning optimized results. ZQV2 denotes the ZeroQuant-V2

## End-to-end Performance and Analysis

Based on Table 4, we can find that Agile-Quant can bring an overall acceleration ratio of 2.3x to 2.6x depending on the model. Especially, Agile-Quant-8 quantizes the activation into INT8 precision and can achieve approximately 1.8x to 1.9x acceleration compared to FP16 in GeMM. Also, combined with the 4-bit compression and concatenation technique, Agile-Quant-4 can further improve this advantage, achieving approximately 1.75x acceleration compared to INT8 multiplication.

## Conclusions and Limitations

In this paper, we propose Agile-Quant, an activation-guided quantization framework for popular LLMs, and design an end-to-end accelerator on multiple edge devices. We introduce the quantization strategy on model weights and activations, and we import token pruning to optimize quantization. We introduce SIMD-based 4-bit multiplier and efficient TRIP matrix multiplication to achieve the 2.55x speedup on hardware devices. Our next step is to explore lower-bit LLMs and design multiple lower-bit multipliers.

## Acknowledgements

This research is mainly funded by the Army Research Office/Army Research Laboratory via grant W911-NF-20-1-0167 to Northeastern University and is partially supported by the National Science Foundation CCF-1937500 and CNS-1909172.

# Bits	Size (GB)	PPL	Android CPU (s)	Raspberry Pi (s)
OPT-125M				
FP16	0.24	27.65	1.03 1×	143.2 1×
ours*-8	0.04	30.37	0.58 1.7×	80.34 1.7×
ours*-4	0.04	36.95	0.44 2.3×	61.16 2.3×
OPT-1.3B				
FP16	2.50	14.63	5.42 1×	754.25 1×
ours*-8	0.49	14.90	2.98 1.8×	410.58 1.8×
ours*-4	0.49	18.20	2.18 2.5×	296.50 2.5×
OPT-2.7B				
FP16	5.00	12.47	8.28 1×	1150.9 1×
ours*-8	0.94	13.19	4.60 1.8×	625.41 1.8×
ours*-4	0.94	16.32	3.31 2.5×	455.64 2.5×
LLaMA-7B				
FP16	13.5	5.68	10.6 1×	1473.4 1×
ours*-8	2.53	6.09	5.89 1.8×	810.25 1.8×
ours*-4	2.53	8.81	4.44 2.4×	610.98 2.3×

Table 4: Hardware results under different data precision for various LLMs. Results are obtained by Agile-Quant with 4 or 8 bits and token pruning.

## References

- ARM. 2023. A collection of low-level machine learning functions optimized with SIMD technologies. <https://arm-software.github.io/ComputeLibrary/v22.05/>.
- Brown, T.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J. D.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; et al. 2020a. Language models are few-shot learners. *NeurIPS*, 33: 1877–1901.
- Brown, T. B.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; Agarwal, S.; Herbert-Voss, A.; Krueger, G.; Henighan, T.; Child, R.; Ramesh, A.; Ziegler, D. M.; Wu, J.; Winter, C.; Hesse, C.; Chen, M.; Sigler, E.; Litwin, M.; Gray, S.; Chess, B.; Clark, J.; Berner, C.; McCandlish, S.; Radford, A.; Sutskever, I.; and Amodei, D. 2020b. Language Models are Few-Shot Learners.
- Cai, J.; Takemoto, M.; and Nakajo, H. 2018. A Deep Look into Logarithmic Quantization of Model Parameters in Neural Networks. In *IAIT*.
- Dettmers, T.; Lewis, M.; Belkada, Y.; and Zettlemoyer, L. 2022. Llm.int8(): 8-bit matrix multiplication for transformers at scale. *arXiv preprint arXiv:2208.07339*.
- 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*.
- Dong, P.; Sun, M.; Lu, A.; Xie, Y.; Liu, K.; Kong, Z.; Meng, X.; Li, Z.; Lin, X.; Fang, Z.; et al. 2023. Heatvit: Hardware-efficient adaptive token pruning for vision transformers. In *HPCA*, 442–455. IEEE.
- Dukhan, M.; Wu, Y.; and Lu, H. 2018. QNNPACK: Open source library for optimized mobile deep learning.
- Frantar, E.; Ashkboos, S.; Hoefler, T.; and Alistarh, D. 2022. GPTQ: Accurate Post-training Compression for Generative Pretrained Transformers. *arXiv*.
- Jacob, B.; Kligys, S.; Chen, B.; Zhu, M.; Tang, M.; Howard, A.; Adam, H.; and Kalenichenko, D. 2018. Quantization and training of neural networks for efficient integer-arithmetic-only inference. In *CVPR*, 2704–2713.
- Jacob, B.; and Warden, P. 2017. gemmlowp: A small self-contained low-precision gemm library. *Retrieved June*, 14: 2018.
- Kim, S.; Shen, S.; Thorsley, D.; Gholami, A.; Kwon, W.; Hassoun, J.; and Keutzer, K. 2022. Learned token pruning for transformers. In *KDD*, 784–794.
- Kong, Z.; Ma, H.; Yuan, G.; Sun, M.; Xie, Y.; Dong, P.; Meng, X.; Shen, X.; Tang, H.; Qin, M.; et al. 2022. Peeling the Onion: Hierarchical Reduction of Data Redundancy for Efficient Vision Transformer Training. *arXiv*.
- Liang, Y.; Ge, C.; Tong, Z.; Song, Y.; Wang, J.; and Xie, P. 2022. Not all patches are what you need: Expediting vision transformers via token reorganizations. *arXiv*.
- Lin, J.; Tang, J.; Tang, H.; Yang, S.; Dang, X.; and Han, S. 2023. AWQ: Activation-aware Weight Quantization for LLM Compression and Acceleration. *arXiv*.
- Lin, Y.; Zhang, T.; Sun, P.; Li, Z.; and Zhou, S. 2022. FQ-ViT: Post-Training Quantization for Fully Quantized Vision Transformer. In *IJCAI*, 1173–1179.
- Merity, S.; Xiong, C.; Bradbury, J.; and Socher, R. 2016. Pointer sentinel mixture models. *arXiv*.
- Radford, A.; Wu, J.; Child, R.; Luan, D.; Amodei, D.; Sutskever, I.; et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8): 9.
- Raffel, C.; Shazeer, N.; Roberts, A.; Lee, K.; Narang, S.; Matena, M.; Zhou, Y.; Li, W.; and Liu, P. J. 2020. Exploring the Limits of Transfer Learning with a Unified Text-to-Text Transformer. *Journal of Machine Learning Research*.
- Scao, T. L.; Fan, A.; Akiki, C.; Pavlick, E.; Ilić, S.; Hesslow, D.; Castagné, R.; Luccioni, A. S.; Yvon, F.; Gallé, M.; et al. 2022. Bloom: A 176b-parameter open-access multilingual language model. *arXiv*.
- Shen, X.; Kong, Z.; Qin, M.; Dong, P.; Yuan, G.; Meng, X.; Tang, H.; Ma, X.; and Wang, Y. 2023. Data Level Lottery Ticket Hypothesis for Vision Transformers. In *IJCAI*.
- Touvron, H.; Lavril, T.; Izacard, G.; Martinet, X.; Lachaux, M.-A.; Lacroix, T.; Rozière, B.; Goyal, N.; Hambro, E.; Azhar, F.; Rodriguez, A.; Joulin, A.; Grave, E.; and Lample, G. 2023. LLaMA: Open and Efficient Foundation Language Models. *arXiv*.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. *NeurIPS*, 30.
- Wu, X.; Yao, Z.; and He, Y. 2023. ZeroQuant-FP: A Leap Forward in LLMs Post-Training W4A8 Quantization Using Floating-Point Formats. *arXiv preprint arXiv:2307.09782*.
- Xiao, G.; Lin, J.; Seznec, M.; Wu, H.; Demouth, J.; and Han, S. 2022. SmoothQuant: Accurate and Efficient Post-Training Quantization for Large Language Models. *arXiv*.
- Zhang, S.; Roller, S.; Goyal, N.; Artetxe, M.; Chen, M.; Chen, S.; Dewan, C.; Diab, M.; Li, X.; Lin, X. V.; et al. 2022. Opt: Open pre-trained transformer language models. *arXiv*.
- Zhang, Y.; Zhao, L.; Cao, S.; Wang, W.; Cao, T.; Yang, F.; Yang, M.; Zhang, S.; and Xu, N. 2023. Integer or Floating Point? New Outlooks for Low-Bit Quantization on Large Language Models. *arXiv*.