

PTQ4SAM: Post-Training Quantization for Segment Anything

Chengtao Lv^{1*}, Hong Chen^{1*}, Jinyang Guo^{1,3†}, Yifu Ding¹, Xianglong Liu^{1,2,4}

¹ State Key Laboratory of Complex & Critical Software Environment, Beihang University

² Zhongguancun Laboratory ³ Institute of Artificial Intelligence, Beihang University

⁴ Institute of data space, Hefei Comprehensive National Science Center

Abstract

*Segment Anything Model (SAM) has achieved impressive performance in many computer vision tasks. However, as a large-scale model, the immense memory and computation costs hinder its practical deployment. In this paper, we propose a post-training quantization (PTQ) framework for Segment Anything Model, namely PTQ4SAM. First, we investigate the inherent bottleneck of SAM quantization attributed to the bimodal distribution in *post-Key-Linear* activations. We analyze its characteristics from both *per-tensor* and *per-channel* perspectives, and propose a *Bimodal Integration* strategy, which utilizes a mathematically equivalent sign operation to transform the bimodal distribution into a relatively easy-quantized normal distribution offline. Second, SAM encompasses diverse attention mechanisms (*i.e.*, self-attention and two-way cross-attention), resulting in substantial variations in the *post-Softmax* distributions. Therefore, we introduce an *Adaptive Granularity Quantization for Softmax* through searching the optimal power-of-two base, which is hardware-friendly. Extensive experimental results across various vision tasks (instance segmentation, semantic segmentation and object detection), datasets and model variants show the superiority of PTQ4SAM. For example, when quantizing SAM-L to 6-bit, we achieve loss-less accuracy for instance segmentation, about 0.5% drop with theoretical $3.9\times$ acceleration. The code is available at <https://github.com/chengtao-lv/PTQ4SAM>.*

1. Introduction

With remarkable zero-shot ability and user-friendly flexible prompt technique, Segment Anything Model (SAM) [20] has recently become a novel foundation model in a range of generic vision applications, including image segmentation [3, 21, 23, 70], object detection [50, 54], tracking [4, 38, 63]

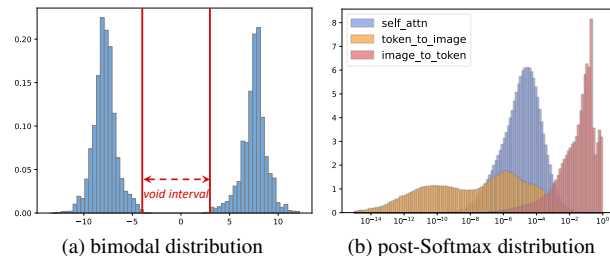


Figure 1. The histogram of two special distributions in SAM: (a) bimodal distribution in *post-Key-Linear* activations. (b) *post-Softmax* distributions of self-attention, token-to-image cross-attention and image-to-token cross-attention.

and other downstream tasks [31, 33, 49, 52, 64, 66]. However, the transformer architectures in SAM require intensive computation and memory footprint, which hurdles the practical deployment on resource-constrained edge-devices.

To address this issue, several quantization approaches [5, 10, 17, 26, 27, 42, 43, 56, 59] were proposed to convert weights and activations from floating-point to low-bit. There are two categories of quantization methods: 1) Quantization-Aware Training (QAT) and 2) Post-Training Quantization (PTQ). QAT retrains a model by utilizing the whole labeled training dataset, which will be time-consuming due to the corresponding massive dataset (SA-1B). On the other hand, PTQ is more promising because it only requires small unlabeled samples to calibrate the pre-trained networks. In this paper, we focus on designing the PTQ approach as it is more effective in practical usage.

Although previous PTQ methods have showcased significant accomplishments in various scenarios, including convolutional neural networks (CNNs) [5, 26, 42, 56, 59], vision transformers (ViTs) [7, 27, 67] as well as large language models (LLMs) [57, 58, 61], directly adopting these methods to Segment Anything Models will raise two unique challenges that necessitate a revisiting of traditional PTQ schemes: 1) we observe the bimodal distribution appears in *post-Key-Linear* activations, *i.e.*, the output activations of key linear, as shown in Figure 1a. The two peaks

* Equal contribution.

† Corresponding author: jinyanguo@buaa.edu.cn

and their central void interval severely enlarge the range of the entire distribution, which negatively affects quantization performance. 2) Owing to the diverse kinds of attention mechanisms, SAMs exhibit more complicated post-Softmax distributions when compared to ViTs, as shown in Figure 1b. Statistically, about 72.5% post-Softmax activations of `image-to-token` are more than 0.01 while only 0.4% in `token-to-image`. Prior works [7, 27, 30, 67] have not adequately addressed this discrepancy and treated them equally, causing the potential loss of inherent information. Therefore, it is desirable to design specialized components for post-Softmax distribution in SAM.

Based on the above observations, in this paper, we propose a novel post-training framework called PTQ4SAM specifically designed for Segment Anything Model quantization. First, we present a *Bimodal Integration* (BIG) strategy to deftly eliminate the bimodal distribution. Specifically, we conduct an in-depth analysis of bimodal distribution from both per-tensor and per-channel perspectives, finding that the bimodal distribution is the crucial obstacle to SAM quantization. We leverage its per-tensor characteristic to determine whether a distribution bimodal distribution, and use its per-channel characteristic to transfer this bimodal distribution to a normal distribution by simultaneously absorbing sign factor into query linear and key linear *offline*. This transformation is mathematically equivalent and significantly narrows the distribution range for better quantization performance. Second, we propose an *Adaptive Granularity Quantization* (AGQ) explicitly tailored for diverse post-Softmax distributions, which affords a suitable trade-off in granularity for both lower and higher attention scores. We provide theoretical proof for its efficiency on hardware by searching the optimal power-of-two base. Instead of minimizing quantization errors of attention scores, we design the objective by the matrix multiplication output between attention scores and values, which is more robust and beneficial to the ultimate performance.

We conduct extensive experiments on fundamental tasks and different model variants to demonstrate the versatility of PTQ4SAM. Our PTQ4SAM can seamlessly plug into both statistic-based and learning-based PTQ methods, achieving 3.9× FLOPs and 4.9× storage savings while maintaining lossless performance on 6-bit SAM-L and SAM-H. Our major contributions are summarized as follows:

- To our best knowledge, our work is the first post-training quantization solution tailored for Segment Anything Model, dubbed PTQ4SAM.
- We observe a challenging bimodal distribution for quantization and analyze its characteristics. To overcome it, we propose a Bimodal Integration (BIG) strategy, which automatically detects it and transforms the bimodal distribution to normal distribution equivalently.
- We present the Adaptive Granularity Quantization (AGQ)

which represents diverse post-Softmax distributions accurately with appropriate granularity.

- Comprehensive experiments conducted on various tasks, variants, and bit-widths demonstrate our PTQ4SAM is a plug-and-play method and significantly surpasses previous state-of-the-art PTQ schemes by a large margin.

2. Related Work

2.1. Segment Anything

Recently, Meta AI Research has revolutionarily approached a general, promptable Segment Anything Model (SAM) [20]. Pre-training on web-scale datasets (SA-1B), SAM demonstrates the capability to generalize across diverse downstream tasks [3, 49, 53, 63, 66]. HQ-SAM [19] designs learnable tokens and global-local fusion schemes to obtain high-quality masks. SEEM [74] extends the referring image to prompt types and integrates a joint visual-semantic space. In the realm of medical research, MedSAM [36] and SAM-Med2D [3] fine-tune SAM through large-scale medical image datasets. Combined with a series of visual-language models [24, 40, 46], Anything-3D [49] and SA3D [1] applies SAM to the single-view 3D reconstruction task while Seal [34] to the 3D point cloud segmentation task. Suffering from substantial computational requirements, some efficient SAMs, including MobileSAM [68] and FastSAM [71] are introduced. However, SAM still undergoes untenable resource-intensive consumption. Its real-time processing capabilities have received widespread expectations.

2.2. Post-Training Quantization

As a predominant compression approach [11–16], the mainstream post-training quantization (PTQ) methods can be broadly divided into two categories [44, 72]: statistic-based PTQ and learning-based PTQ. Statistic-based PTQ methods solely seek optimal quantization parameters to minimize quantization errors, whereas learning-based PTQ methods fine-tune model parameters, including both weights and quantization parameters. Our approach is out-of-the-box on both kinds of methods.

2.2.1 Statistic-Based PTQ

Numerous classic statistic-based quantization methods [5, 17, 39, 41, 43] have been shown to achieve minimal loss in precision primarily for convolutional neural networks. However, with the widespread popularity of networks featuring novel architectures, researchers have introduced quantization schemes specifically designed for these networks. When quantizing ViTs, twin uniform quantization [67], Log-Int-Softmax [30], scale reparameterization [27] and matthew-effect preserving quantization [7] are proposed to tackle the output distribution from soft-

max. The LLM quantization techniques include weight-only quantization [2, 6, 8, 28], weight and activation quantization [57, 58, 61, 65], aiming to settle the outlier issue from the activations. PTQ4DM [48], Q-Diffusion [25] discover the variations in the activations during multiple denoising steps in Diffusion and design specialized calibration strategies. However, there is a notable gap between these distinctive distributions and SAMs, and we are the first to explore quantizing SAMs for efficient inference.

2.2.2 Learning-Based PTQ

Based on statistic-based PTQ methods, several learning-based PTQ schemes were also proposed. AdaRound [42] optimizes the rounding operation when quantizing weights to minimize the overall loss of the model. Subsequently, many methods are proposed based on AdaRound. BRECC [26] proposes a block-wise reconstruction algorithm to optimize the quantized model. QDrop [56] introduces drop operation during the reconstruction process to increase the flatness of the optimized model. PD-Quant [32] introduces global information when optimizing quantization parameters. MRECG [37] focuses on the oscillation problem in PTQ and FlexRound [22] proposes a new learnable weight rounding scheme on large language models for the first time. Unfortunately, these techniques are mainly carried out based on CNN architecture models. Transformer architecture models like SAM remain unexplored.

3. Method

In this section, we will introduce our PTQ4SAM in detail. First, we introduce uniform quantization and logarithmic quantization in Section 3.1. Subsequently, we present an *Bimodal Integration* (BIG) strategy to eliminate bimodal distribution in Section 3.2. Finally, we analyze the discrepancy of diverse post-softmax distribution and propose *Adaptive Granularity Quantization* (AGQ) in Section 3.3.

3.1. Preliminaries

Basic Notations. We use \mathbf{X} to represent a matrix, whereas the vectors are marked by \mathbf{x} . The operator \odot is used to represent element-wise multiplication between matrices or vectors and operator \cdot denotes scalar multiplication. Also, we use \mathbf{XW} to denote matrix multiplication.

Post-training Quantization. Post-training quantization is a prevalent approach to compress the pre-trained neural network. In this paper, we merely study the hardware-efficient quantization methods. For uniform quantization, quantization and de-quantization operations can be defined as:

$$x_q = \text{clamp}(\lfloor \frac{x}{s} \rfloor + z, 0, 2^k - 1), \quad (1)$$

$$\hat{x} = s \cdot (x_q - z) \approx x, \quad (2)$$

where s and z denote the scaling factor and zero point, respectively. $\lfloor \cdot \rfloor$ is the round-to-nearest operator. x and \hat{x} are floating-point and de-quantized values, and x_q is mapped integer. `clamp` function clips the values fall outside the range of a k -bit integer.

In light of rapid bit-shifting operations, Log2 Quantization has emerged as an alternative hardware-oriented quantization approach. Due to the Log2 Quantization is exclusively employed on post-Softmax activations, it is simply formulated as:

$$x_q = \text{clamp}(\lfloor -\log_2 \frac{x}{s} \rfloor, 0, 2^k - 1), \quad (3)$$

$$\hat{x} = s \cdot 2^{-x_q} \approx x. \quad (4)$$

3.2. Bimodal Integration

Intuitively, the bimodal distribution poses significant challenges for quantization. The two peaks, accompanied by their central void or sparse interval considerably expand the distribution range, leading to over $5\times$ quantization errors compared to normal distribution experimentally. Therefore, we first make an in-depth analysis of bimodal distribution in SAM from two perspectives: 1) From **per-tensor** perspective: the distribution contains two peaks and their centers are symmetric, *e.g.*, -8 and 8 in Figure 1a. 2) From **per-channel** perspective: the activations of each channel only persist in a fixed peak, indicating pronounced asymmetry inside one channel. As shown in Figure 3, for instance, the activations of the 0-th channel correspond to a negative peak (*i.e.*, -8), while the activations of 1-th channel belong to a positive peak (*i.e.*, 8). Generally, about half channels (*e.g.*, 46.1% in SAM-B) cluster in the positive peak and the remaining channels cluster in the negative peak.

Recently, some methodologies [27, 58] have been introduced to overcome this channel-wise asymmetry by equivalently adjusting the weights of LayerNorm and the subsequent linear. However, we observe that the bimodal distribution does not exist in `post-LayerNorm` activations but rather prominently concentrates in `post-Key-Linear` activations (Figure 2), rendering the aforementioned methods inapplicable. To precisely estimate the influence of bimodal distribution, the following matrix multiplication between \mathbf{Q} and \mathbf{K} can be formulated as

$$\mathbf{QK}^\top = \underbrace{(\mathbf{X}_q \mathbf{W}_q + \mathbf{b}_q)}_{\text{normal distribution}} \underbrace{(\mathbf{X}_k \mathbf{W}_k + \mathbf{b}_k)}_{\text{bimodal distribution}}^\top, \quad (5)$$

where $\mathbf{W} \in \mathbb{R}^{m \times n}$ and $\mathbf{b} \in \mathbb{R}^n$ are the weight and bias of linear (*i.e.*, fully connected layer). m and n are input and output feature dimensions. The subscripts q and k denote the query and key linears, respectively. Note that the matrix multiplication of \mathbf{Q} and \mathbf{K} essentially represents the matrix multiplication of normal and bimodal distributions. Motivated by the above analysis, we adopt a channel-wise sign

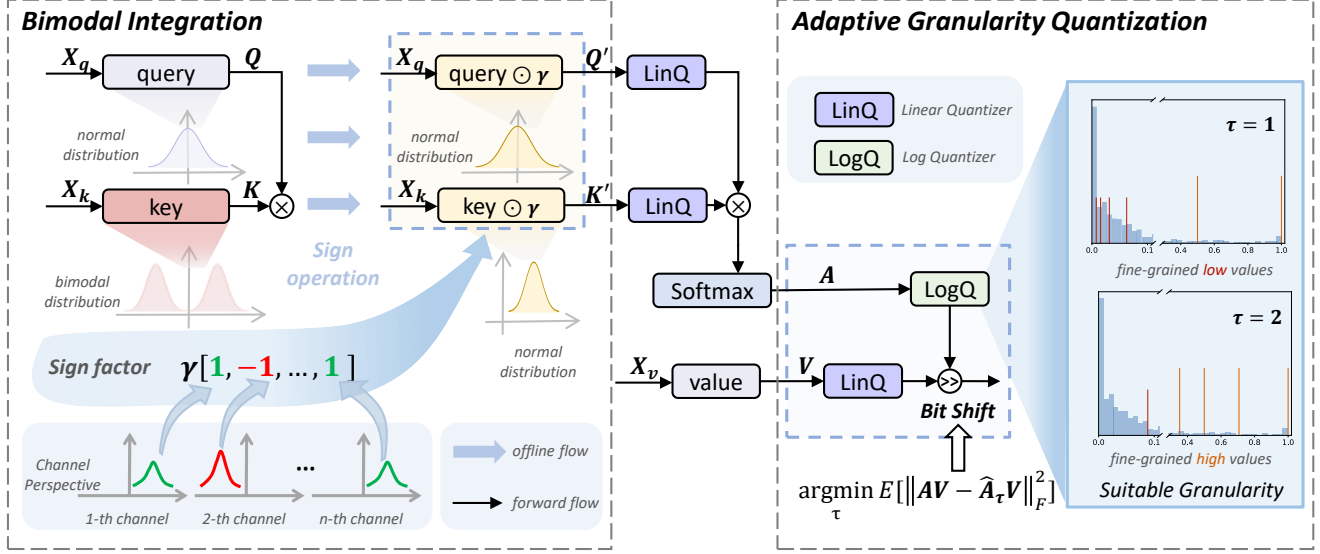


Figure 2. Illustration of our proposed PTQ4SAM. The Bimodal Integration eliminates the bimodal distribution by simultaneously multiplying a channel-wise γ to both the query and key linears. The Adaptive Granularity Quantization is employed for post-softmax distribution.

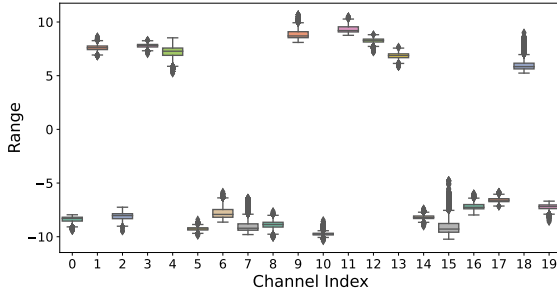


Figure 3. Boxplot of different channels of post-Key-Linear activations in SAM.

factor $\gamma \in \mathbb{R}^n$ to transfer the bimodal distribution to a normal distribution. We assume that γ can be computed from the mean value of each channel, considering the sign factor of j -th channel:

$$\gamma_j = \begin{cases} +1, & \text{if } \text{mean}(\mathbf{K}_{:,j}) \geq 0 \\ -1, & \text{otherwise} \end{cases}, \quad (6)$$

where $\mathbf{K}_{:,j}$ denotes the post-Key-Linear activations in j -th channel. And we compute the γ_j through its sign of mean value. Specifically, if the j -th channel is negative peak, γ_j will be -1. Then we multiply it to the corresponding channel of query linear and key linear simultaneously to maintain equivalence. After that, $\mathbf{K}_{:,j}$ will be transferred to the positive. Conversely, if one channel is positive peak, it will remain invariant as its sign factor is 1. As shown in Figure 2, after using our BIG strategy, the negative peak in \mathbf{K} has been merged into the positive peak, thereby transferring the bimodal distribution to a normal distribution. On

the other hand, \mathbf{Q} still preserve normal distribution:

$$\begin{aligned} \mathbf{Q}\mathbf{K}^\top &= ((\mathbf{X}_q \mathbf{W}_q + \mathbf{b}_q) \odot \gamma) ((\mathbf{X}_k \mathbf{W}_k + \mathbf{b}_k)^\top \odot \gamma^\top) \\ &= \underbrace{(\mathbf{X}_q \mathbf{W}'_q + \mathbf{b}'_q)}_{\text{normal distribution}} \underbrace{(\mathbf{X}_k \mathbf{W}'_k + \mathbf{b}'_k)^\top}_{\text{normal distribution}}. \end{aligned} \quad (7)$$

Note the sign factor γ can be easily absorbed into previous query linear and key linear offline without any computation overhead, *i.e.*, $\mathbf{W}' = \mathbf{W} \odot \gamma$ and $\mathbf{b}' = \mathbf{b} \odot \gamma$.

bimodal discovery: However, not all post-Key-Linear activations in SAM is bimodal distribution. To discriminate the bimodal distribution, we first adopt the Gaussian kernel density estimation to compute the probability density function (PDF) [47] by the whole tensor. Based on the continuous and smooth function, we quantitatively describe the peaks as local maxima. To avoid recognizing two small bumps as two peaks, we constrain the peak height and the distance between two peaks. More implementation details are described in supplementary materials.

In summary, our *Bimodal Integration* (BIG) strategy comprises three steps: bimodal discovery, γ computation and equivalent transformation. Due to the strong asymmetry, only one sample is enough to compute the sign factor γ , described in Algorithm 1. Therefore, our BIG is efficient and the extra computational burden can be ignored.

3.3. Adaptive Granularity Quantization

In transformer architecture, the softmax function is utilized to convert attention scores into probabilities, ensuring that the scores are normalized and lie between 0 and 1.

Algorithm 1 PTQ4SAM CALIBRATION

Input: Full-precision SAM M_F , Calibration Set $Calib$
Output: Quantized model M_Q

```

1: for  $i = 1 : \#Calib$  do
2:    $\triangleright$  Bimodal Integration
3:   if  $i == 1$  then
4:     Discriminate bimodal distribution
5:     Compute sign factor  $\gamma$  with Eq. 6
6:     Equivalent transformation with Eq. 7
7:   end if
8:   for  $l = 1 : \#Quantizer$  do
9:     Initialize  $l$ -th quantizer's scaling factor
10:     $\triangleright$  For post-Softmax Distribution
11:    if Adaptive Granularity Quantizer then
12:      Compute and add error of each  $\tau$ 
13:      if  $i == \#Calib$  then
14:        Determine optimal  $\tau$  with Eq. 14
15:      end if
16:    end if
17:  end for
18: end for
19: return  $M_Q$ 

```

[7, 27, 30, 67] identify the extremely unbalanced power-law distribution as the rationale of quantization difficulty and devise a specialized quantizer to address it. However, the aforementioned methods are mainly designed for the softmax in self-attention mechanism. SAM also incorporates cross-attention in two directions, *i.e.*, `token-to-image` and `image-to-token` cross-attention, amplifying the remarkable discrepancy between post-softmax distributions in Figure 1b. For example, there are more ultra-low-values in `token-to-image`, displaying a smooth distribution under a logarithmic scale. On the contrary, the distributions in `image-to-token` and `self-attention` exhibit higher kurtosis and exist more high values.

To tackle this discrepancy, we revisit the logarithmic quantizer and propose an *Adaptive Granularity Quantization* (AGQ) with an adaptive parameter τ to adjust the base. As shown in Figure 2, a smaller τ can represent lower attention scores. And as τ increases, the higher attention scores become more fine-grained. Our AGQ, equipped with a suitable τ , achieves a flexible trade-off between the granularity of low and high values under diverse post-Softmax scenarios and different bit-widths. The corresponding quantization and de-quantization operations of AGQ for attention score a (one element in the attention map for simplicity) can be rewritten as:

$$a_q = \text{clamp}(\lfloor -\log_{2^{\frac{1}{\tau}}} \frac{a}{s} \rfloor, 0, 2^k - 1), \quad (8)$$

$$\hat{a} = s_a \cdot 2^{-\frac{a_q}{\tau}}. \quad (9)$$

In our implementation, we use $\tau \in \{2^0, 2^1, \dots, 2^n\}$ for the hardware efficiency, which will be demonstrated later.

The premise for executing the bit-shifting operation is the assurance that the exponential term is an integer. Regrettably, $\frac{a_q}{\tau}$ is not necessarily an integer. Consequently, we first decompose into its integral and fractional components:

$$\hat{a} = s_a \cdot 2^{\lfloor -\frac{a_q}{\tau} \rfloor} \cdot 2^{\frac{(-a_q)\% \tau}{\tau}}, \quad (10)$$

where $\%$ denotes the modulo function and $\lfloor \cdot \rfloor$ denotes the floor function. Subsequently, the multiplication between quantized attention score \hat{a} and value \hat{v} are written as:

$$\hat{a} \cdot \hat{v} = s_a \cdot s_v \cdot 2^{\frac{(-a_q)\% \tau}{\tau}} \cdot 2^{\lfloor -\frac{a_q}{\tau} \rfloor} \cdot v_q \quad (11)$$

$$= s_a \cdot s_v \cdot \underbrace{2^{\frac{(-a_q)\% \tau}{\tau}}}_{(12-1)} \cdot v_q \gg \lceil \frac{a_q}{\tau} \rceil, \quad (12)$$

where \gg indicates the bit-shifting operations. As analyzed above, the term (12-1) is an uncertain floating-point value, it limits the efficient integer-only arithmetic in hardware. However, the term (12-1) possesses only τ discrete values so a small lookup table can be used to avoid it:

$$\hat{a} \cdot \hat{v} = s_a \cdot s_v \cdot LUT((-a_q)\% \tau, v_q \gg \lceil \frac{a_q}{\tau} \rceil), \quad (13)$$

where LUT denotes the small lookup table, which is available on various Neural Network Accelerators, *e.g.* on FPGA. The entries in the LUT are only determined by $(-a_q)\% \tau$ and $v_q \gg \lceil \frac{a_q}{\tau} \rceil$. The first term can be represented with an n -bit number, and the second term with a k -bit (activation bit-width). Notably, since the LUTs for $\tau \in \{2^0, 2^1, \dots, 2^{n-1}\}$ can be incorporated into the LUT for $\tau = 2^n$, the entire network only requires a single LUT. Considering a scenario with 8-bit activation and maximum τ from the set (2^2 for implementation), the size of the LUT is computed as $2^{8+2} \times 4$ bytes = 4KB, which is negligible compared to quantized SAMs.

With the theoretical validation established, our aim is to define an objective to select the optimal τ . To this end, a natural choice is to minimize the local quantization error of attention map \mathbf{A} directly. However, we discover it is inconsistent with the quantization error of its associated attention block, which induces instabilities in the global quantization performance, especially at low-bit (see more details in supplementary materials). Therefore, to alleviate this inconsistency, we design the objective function to measure the quantization error of the matrix multiplication output between attention map \mathbf{A} and values \mathbf{V} :

$$\arg \min_{\tau} \mathbb{E}[\|\mathbf{A}\mathbf{V} - \hat{\mathbf{A}}_{\tau}\mathbf{V}\|_F^2], \quad (14)$$

where $\hat{\mathbf{A}}_{\tau}$ represents the quantized attention map by our AGQ with the base $2^{\frac{1}{\tau}}$ and $\|\cdot\|_F^2$ denotes the Frobenius norm to measure the difference. As elaborated in Algorithm 1, we sort total quantization errors of each τ across the calibration set, then choose the optimal τ after calibration.

Detector	Methods	SAM-B			SAM-L			SAM-H		
		FP	W6A6	W4A4	FP	W6A6	W4A4	FP	W6A6	W4A4
Faster R-CNN [45]	MinMax [17]	33.4	9.2	-	36.4	32.9	-	37.2	31.9	-
	Percentile [59]		10.9	-		33.5	-		32.0	-
	OMSE [5]		11.9	-		33.9	5.4		33.1	7.4
	PTQ4SAM-S		15.4	-		35.7	18.1		36.0	24.1
	AdaRound [42]		23.1	-		34.3	8.7		33.7	14.5
	BRECQ [26]		24.1	-		34.2	10.7		33.7	15.1
	QDrop [56]		29.3	13.0		35.2	22.6		36.3	32.3
PTQ4SAM-L	30.3	16.0	35.8	28.7	36.5	33.5				
YOLOX [9]	MinMax [17]	37.0	10.7	-	40.4	37.5	-	41.0	36.1	-
	Percentile [59]		12.0	-		38.0	-		36.3	-
	OMSE [5]		13.5	-		38.4	6.1		37.5	7.8
	PTQ4SAM-S		17.4	-		40.0	20.6		40.3	26.7
	AdaRound [42]		26.4	-		38.9	11.1		38.3	16.7
	BRECQ [26]		26.1	-		38.9	12.0		38.3	16.3
	QDrop [56]		33.6	13.3		39.7	25.3		40.4	35.8
PTQ4SAM-L	34.3	18.4	40.3	31.6	40.7	37.6				
H-Deformable-DETR [18]	MinMax [17]	38.2	10.9	-	41.5	38.6	-	42.0	37.3	-
	Percentile [59]		12.3	-		39.0	-		37.5	-
	OMSE [5]		15.0	-		39.6	6.2		38.6	7.7
	PTQ4SAM-S		17.9	-		41.0	20.9		41.3	27.3
	AdaRound [42]		27.2	-		39.9	8.0		39.4	16.3
	BRECQ [26]		27.9	-		39.9	11.1		39.5	15.5
	QDrop [56]		34.3	13.2		40.5	25.8		41.4	36.5
PTQ4SAM-L	35.1	17.3	41.2	32.1	41.6	38.4				
DINO [69]	MinMax [17]	44.5	11.2	-	48.6	44.7	-	49.1	42.8	-
	Percentile [59]		14.0	-		45.4	-		43.1	-
	OMSE [5]		16.6	-		45.9	6.8		44.5	8.3
	PTQ4SAM-S		20.4	-		47.7	23.1		48.1	30.5
	AdaRound [42]		31.2	1.2		46.6	8.8		46.0	18.2
	BRECQ [26]		31.8	3.6		46.6	12.3		46.0	17.6
	QDrop [56]		38.9	11.2		47.5	27.5		48.3	41.7
PTQ4SAM-L	40.4	14.4	48.3	36.6	48.7	43.9				

Table 1. Quantization results of instance segmentation on COCO dataset among different detectors. We provide PTQ4SAM-S and PTQ4SAM-L for both statistic-based and learning-based version. - indicates the final result (mAP) is below 1.

4. Experiments

4.1. Experimental Setup

Tasks, datasets and metrics. We conduct experiments on three mainstream vision tasks. For the instance segmentation task, we utilize predicted boxes generated by the detector as box prompts for SAM to gain accurate binary masks and evaluate its effectiveness on MS-COCO [29] dataset with the metric mean Average Precision (mAP). For semantic segmentation task, the overall framework comprises two branches, and we leverage the fine-grained mask produced by SAM to refine the blurry and imprecise mask boundaries generated by the original segmentor. We evaluate its effectiveness on ADE20K [73] dataset using the mean Intersection over Union (mIOU) as the performance metric. For oriented object detection task, we obtain the final rotated RBoxes by the minimum circumscribed rectangle operation on the masks generated by SAM. Our evaluation of its effectiveness on the DOTA-v1.0 [60] dataset uses mAP.

Implementation details. We choose CNN-based Faster R-CNN [45], YOLOX [9], FCOS [51] and transformer-based H-Deformable-DETR [18], DINO [69] as detectors and advanced SegFormer [62] as segmentor. The adaptive parameter τ is searched from the set $\{2^0, 2^1, 2^2\}$. We set the box threshold to 0.05 for CNN-based detectors and affix a set of 100 adaptive anchors for transformer-based detectors. We randomly sample 32 training images as calibration set and only the first sample is utilized for the determination of the bimodal distribution. For a fair comparison, we adopt per-channel asymmetric quantization for weights and per-tensor asymmetric quantization for activations [32, 56]. Following the common settings [56, 67], the first and last layer/block are not quantized. To verify the effectiveness of PTQ4SAM in two kinds of PTQ methods, we integrate our method into statistic-based OMSE [5] and learning-based QDrop [56], called PTQ4SAM-S and PTQ4SAM-L, respectively. For learning-based methods, we adopt MinMax calibration strategy and design attention block, MLP block for block-

Base	Model	Methods	FP	W6A6	W4A4
31.78	+SAM-B	AdaRound [42]	33.15	32.34	31.78
		BRECQ [26]		32.27	31.78
		QDrop [56]		32.57	31.79
		PTQ4SAM-L		32.65	31.85
	+SAM-L	AdaRound [42]	33.61	32.99	31.97
		BRECQ [26]		33.04	31.98
		QDrop [56]		33.58	32.67
		PTQ4SAM-L		33.66	32.82
	+SAM-H	AdaRound [42]	33.63	33.49	32.17
		BRECQ [26]		33.46	32.12
		QDrop [56]		33.49	32.93
		PTQ4SAM-L		33.66	33.10

Table 2. Quantization results of semantic segmentation. SAM refines the outcomes produced by the original semantic segmentor.

wise reconstruction with 20000 iterations.

4.2. Instance Segmentation Results

We compare our PTQ4SAM with statistic-based methods such as MinMax [17], Percentile [59], OMSE [5] and learning-based methods such as AdaRound [42], BRECQ [26], QDrop [56]. Table 1 lists the performance of all methods. Our method consistently outperforms other methods by a large margin among different detectors. Our 4-bit PTQ4SAM-S achieves comparable results, and exceeds the baseline OMSE by over 10% mAP (*e.g.*, from 5.4% to 18.1% with Faster R-CNN) on SAM-L and about 20% mAP (*e.g.*, from 8.3% to 30.5% with DINO) on SAM-H, recovering them to a usable level. Our PTQ4SAM-S even remarkably surpasses the state-of-the-art learning-based method QDrop, at W6A6 on SAM-L. Meanwhile, our PTQ4SAM-L encouragingly achieves lossless accuracy. For instance, at W6A6 setting, our PTQ4SAM-L achieves 40.3% and 41.2% when applying YOLOX and H-Deformable-DETR on SAM-L, with only 0.1% and 0.3% performance drop compared to full-precision models. When quantizing to more challenging case W4A4, AdaRound and BRECQ become infeasible while our method surpasses the QDrop by 5.1% on SAM-B with YOLOX and 6.3% on SAM-L with H-Deformable-DETR. Compared with SAM-B and SAM-L, SAM-H exhibits greater robustness when introducing quantization noise, but our PTQ4SAM-L still provides about 2% improvement at W4A4. Applying the state-of-art detector DINO [69], our 6-bit PTQ4SAM-L yields 40.4% mAP on SAM-B and achieves a mAP nearing 50% on SAM-L and SAM-H.

4.3. Semantic Segmentation Results

In this part, we verify the performance of the learning-based PTQ methods for semantic segmentation task. The original semantic segmentor achieves 31.78% on ADE20K dataset, and SAMs adjust the preliminary outcomes of seg-

Model	Methods	FP	W6A6	W4A4
SAM-B	AdaRound [42]	64.1	34.05	-
	BRECQ [26]		34.40	-
	QDrop [56]		59.27	41.96
	PTQ4SAM-L		60.33	44.18
SAM-L	AdaRound [42]	64.2	63.44	23.18
	BRECQ [26]		63.60	26.89
	QDrop [56]		63.86	50.11
	PTQ4SAM-L		63.91	56.29
SAM-H	AdaRound [42]	64.6	62.73	24.45
	BRECQ [26]		62.58	25.98
	QDrop [56]		62.83	55.87
	PTQ4SAM-L		64.36	56.01

Table 3. Quantization results of oriented object detection.

mentor, *i.e.*, bring 1.37%, 1.83%, 1.85% enhancement for full-precision SAM-B, SAM-L and SAM-H. As shown in Table 2, quantized SAMs generally still contribute to the final masks. Our method retains the capability of SAM to the greatest extent. In particular, we are surprised to find that 6-bit PTQ4SAM-L even achieves better performance than the full-precision models on both SAM-L and SAM-H. At W4A4 setting, our method provides 1.04% accuracy promotion on SAM-L, outperforming QDrop by 0.15%.

4.4. Object Detection Results

To further demonstrate versatility across other tasks, we test PTQ4SAM-L in oriented object detection. Notably, as shown in Table 3, our method consistently performs better than other learning-based PTQ methods. For instance, when quantizing the network to W6A6, experiments indicate PTQ4SAM-L slightly drops about 0.3% compared with the full-precision model on SAM-L and SAM-H. At the most challenging W4A4 bit-width, AdaRound and BRECQ suffer from non-trivial performance degradation. Contrastively, our method can still obtain satisfactory performance. We achieve over 44% and 56% accuracy on SAM-B and SAM-L, surpassing the baseline QDrop 2.2% and 6.2%.

4.5. Ablation Studies

Ablation for components: Table 4 lists the results of different components. We demonstrate that each component contributes to PTQ4SAM, while the best performance is achieved when both components are jointly applied. Specifically, at relatively higher bit-widths, like W6A6 setting, both BIG and AGQ strategies can bring performance improvement, making the quantized model comparable to the full-precision one. When quantizing SAM to lower bit-widths, *i.e.*, W4A4, BIG can significantly improve the performance, which is 3.4% higher than baseline, as it can minimize the more serious quantization perturbation.

Ablation for different quantizers: We also report the results of different quantizers in Table 5. Notably, simply em-

Row ID	Model	BIG	AGQ	FP	W6A6	W4A4
1	SAM-L	×	×	41.5	40.5	25.8
2		✓	×		40.6	29.2
3		×	✓		41.2	27.3
4		✓	✓		41.2	32.1

Table 4. Ablation study for key components.

#bits	Quantizer	SAM-B	SAM-L	SAM-H
Full-precision	-	37.0	40.4	41.0
W6A6	Uniform	33.6	39.7	40.4
	Log2	33.3	40.2	40.6
	AGQ (ours)	33.9	40.3	40.6
W4A4	Uniform	13.3	25.3	35.8
	Log2	14.1	26.5	37.3
	AGQ (ours)	15.0	27.8	37.6

Table 5. Ablation study for different post-Softmax quantizers.

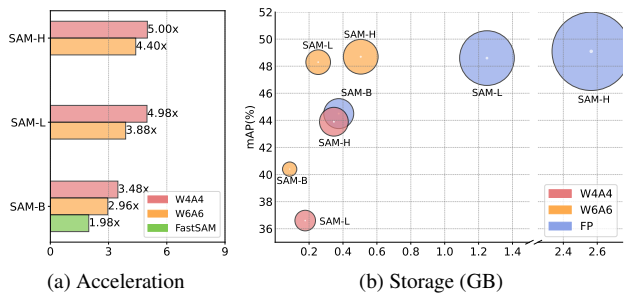


Figure 4. (a) Theoretical acceleration rate (100 prompts) vs. all SAM models. (b) Accuracy vs. storage.

employing Log2 quantizer [41] for attention scores is unstable under different settings (e.g., lower result at W6A6 for SAM-B compared with uniform quantizer). Our AGQ, by contrast, surpasses the uniform quantizer and Log2 quantizer in different cases, boosting 3.5% for 4-bit SAM-L. Apart from the encouraging performance, our AGQ is available on various hardware, ensuring efficient execution.

4.6. Storage Saving and Speedup

We separately calculate the computational complexity in Figure 4a and the memory usage in Figure 4b. Following [35, 55], we use FLOPs to measure the theoretical acceleration effect, which is equal to the amount of 32-bit floating point multiplication plus $\{\frac{1}{8}, \frac{6}{32}\}$ of the amount of $\{4,6\}$ -bit multiplications. Surprisingly, our 6-bit SAM-B (2.96 \times) achieves better acceleration than FastSAM [71] (1.98 \times) while maintaining a close performance. At W4A4, our method reduces computational FLOPs by over 70% and storage by over 85%. As models scale up, both acceleration ratio and memory savings become more significant, while the performance drop becomes less.




Figure 5. Visualization of instance segmentation on 4-bit SAM-L.

4.7. Qualitative Results

To exhibit the superiority of our PTQ4SAM, especially on low-bit quantization (W4A4), we visualize the instance segmentation results in Figure 5 compared with other existing SOTA methods on COCO dataset. We can perceive that most methods fail to produce clear boundaries and miss the salient pixels in the center. For example, other methods do not distinguish the complex edge, even erroneously categorizing the sky as foreground (train on row 1). Besides, these schemes are unable to detect the objects against complicated backgrounds (surfboard on row 3). Furthermore, our method outperforms other methods on ensuring the integrity of the object (toilet and person on rows 2, 4).

5. Conclusion

In this paper we first propose a novel post-training quantization framework, PTQ4SAM for Segment Anything Model. To begin with, we observe the significant bottleneck lies in bimodal distribution and explore its occurrence position. We introduce Bimodal Integration strategy to eliminate the negative effect of quantization caused by such bimodal distribution. We also present the Adaptive Granularity Quantization which can fit diverse post-Softmax distribution by searching the hardware-friendly base. Extensive experiments demonstrate that our method consistently yields gratifying results across various tasks. Nevertheless, the reason for bimodal distribution in SAM remains unclear. This direction serves as a potential avenue for our future research.

References

- [1] Jiazhong Cen, Zanwei Zhou, Jiemin Fang, Wei Shen, Lingxi Xie, Xiaopeng Zhang, and Qi Tian. Segment anything in 3d with nerfs. *arXiv preprint arXiv:2304.12308*, 2023. **2**
- [2] Hong Chen, Chengtao Lv, Liang Ding, Haotong Qin, Xiabin Zhou, Yifu Ding, Xuebo Liu, Min Zhang, Jinyang Guo, Xianglong Liu, et al. Db-llm: Accurate dual-binarization for efficient llms. *arXiv preprint arXiv:2402.11960*, 2024. **3**
- [3] Junlong Cheng, Jin Ye, Zhongying Deng, Jianpin Chen, Tianbin Li, Haoyu Wang, Yanzhou Su, Ziyang Huang, Jilong Chen, Lei Jiang, et al. Sam-med2d. *arXiv preprint arXiv:2308.16184*, 2023. **1, 2**
- [4] Yangming Cheng, Liulei Li, Yuanyou Xu, Xiaodi Li, Zongxin Yang, Wenguan Wang, and Yi Yang. Segment and track anything. *arXiv preprint arXiv:2305.06558*, 2023. **1**
- [5] Yoni Choukroun, Eli Kravchik, Fan Yang, and Pavel Kisilev. Low-bit quantization of neural networks for efficient inference. In *2019 IEEE/CVF International Conference on Computer Vision Workshop (ICCVW)*, pages 3009–3018. IEEE, 2019. **1, 2, 6, 7**
- [6] Tim Dettmers, Mike Lewis, Younes Belkada, and Luke Zettlemoyer. Llm.int8(): 8-bit matrix multiplication for transformers at scale. *arXiv preprint arXiv:2208.07339*, 2022. **3**
- [7] Yifu Ding, Haotong Qin, Qinghua Yan, Zhenhua Chai, Junjie Liu, Xiaolin Wei, and Xianglong Liu. Towards accurate post-training quantization for vision transformer. In *Proceedings of the 30th ACM International Conference on Multimedia*, pages 5380–5388, 2022. **1, 2, 5**
- [8] Elias Frantar, Saleh Ashkboos, Torsten Hoefler, and Dan Alistarh. Gptq: Accurate post-training quantization for generative pre-trained transformers. *arXiv preprint arXiv:2210.17323*, 2022. **3**
- [9] Zheng Ge, Songtao Liu, Feng Wang, Zeming Li, and Jian Sun. Yolox: Exceeding yolo series in 2021. *arXiv preprint arXiv:2107.08430*, 2021. **6**
- [10] Amir Gholami, Sehoon Kim, Zhen Dong, Zhewei Yao, Michael W Mahoney, and Kurt Keutzer. A survey of quantization methods for efficient neural network inference. In *Low-Power Computer Vision*, pages 291–326. Chapman and Hall/CRC, 2022. **1**
- [11] Jinyang Guo, Wanli Ouyang, and Dong Xu. Multi-dimensional pruning: A unified framework for model compression. In *CVPR*, 2020. **2**
- [12] Jinyang Guo, Weichen Zhang, Wanli Ouyang, and Dong Xu. Model compression using progressive channel pruning. *IEEE Transactions on Circuits and Systems for Video Technology*, 2020.
- [13] Jinyang Guo, Jiaheng Liu, and Dong Xu. Jointpruning: Pruning networks along multiple dimensions for efficient point cloud processing. *IEEE Transactions on Circuits and Systems for Video Technology*, 2021.
- [14] Jinyang Guo, Dong Xu, and Guo Lu. Cbanet: Towards complexity and bitrate adaptive deep image compression using a single network. *IEEE Transactions on Image Processing*, 2023.
- [15] Jinyang Guo, Dong Xu, and Wanli Ouyang. Multidimensional pruning and its extension: A unified framework for model compression. *IEEE Transactions on Neural Networks and Learning Systems*, 2023.
- [16] J. Guo, W. Ouyang, and D. Xu. Channel pruning guided by classification loss and feature importance. In *AAAI*, 2020. **2**
- [17] Benoit Jacob, Skirmantas Kligys, Bo Chen, Menglong Zhu, Matthew Tang, Andrew Howard, Hartwig Adam, and Dmitry Kalenichenko. Quantization and training of neural networks for efficient integer-arithmetic-only inference. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2704–2713, 2018. **1, 2, 6, 7**
- [18] Ding Jia, Yuhui Yuan, Haodi He, Xiaopei Wu, Haojun Yu, Weihong Lin, Lei Sun, Chao Zhang, and Han Hu. Detsr with hybrid matching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19702–19712, 2023. **6**
- [19] Lei Ke, Mingqiao Ye, Martin Danelljan, Yifan Liu, Yu-Wing Tai, Chi-Keung Tang, and Fisher Yu. Segment anything in high quality. *arXiv preprint arXiv:2306.01567*, 2023. **2**
- [20] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. Segment anything. *arXiv preprint arXiv:2304.02643*, 2023. **1, 2**
- [21] Xing Lan, Jiayi Lyu, Hanyu Jiang, Kun Dong, Zehai Niu, Yi Zhang, and Jian Xue. Foodsam: Any food segmentation. *arXiv preprint arXiv:2308.05938*, 2023. **1**
- [22] Jung Hyun Lee, Jeonghoon Kim, Se Jung Kwon, and Dongsoo Lee. Flexround: Learnable rounding based on element-wise division for post-training quantization. *arXiv preprint arXiv:2306.00317*, 2023. **3**
- [23] Feng Li, Hao Zhang, Peize Sun, Xueyan Zou, Shilong Liu, Jianwei Yang, Chunyuan Li, Lei Zhang, and Jianfeng Gao. Semantic-sam: Segment and recognize anything at any granularity. *arXiv preprint arXiv:2307.04767*, 2023. **1**
- [24] Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *International Conference on Machine Learning*, pages 12888–12900. PMLR, 2022. **2**
- [25] Xiuyu Li, Yijiang Liu, Long Lian, Huanrui Yang, Zhen Dong, Daniel Kang, Shanghang Zhang, and Kurt Keutzer. Q-diffusion: Quantizing diffusion models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 17535–17545, 2023. **3**
- [26] Yuhang Li, Ruihao Gong, Xu Tan, Yang Yang, Peng Hu, Qi Zhang, Fengwei Yu, Wei Wang, and Shi Gu. Brecq: Pushing the limit of post-training quantization by block reconstruction. *arXiv preprint arXiv:2102.05426*, 2021. **1, 3, 6, 7**
- [27] Zhikai Li, Junrui Xiao, Lianwei Yang, and Qingyi Gu. Repqvit: Scale reparameterization for post-training quantization of vision transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 17227–17236, 2023. **1, 2, 3, 5**
- [28] Ji Lin, Jiaming Tang, Haotian Tang, Shang Yang, Xingyu Dang, and Song Han. Awq: Activation-aware weight quantization for llm compression and acceleration. *arXiv preprint arXiv:2306.00978*, 2023. **3**

- [29] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part V 13*, pages 740–755. Springer, 2014. 6
- [30] Yang Lin, Tianyu Zhang, Peiqin Sun, Zheng Li, and Shuchang Zhou. Fq-vit: Post-training quantization for fully quantized vision transformer. *arXiv preprint arXiv:2111.13824*, 2021. 2, 5
- [31] Aishan Liu, Xianglong Liu, Jiaxin Fan, Yuqing Ma, Anlan Zhang, Huiyuan Xie, and Dacheng Tao. Perceptual-sensitive gan for generating adversarial patches. In *AAAI*, 2019. 1
- [32] Jiawei Liu, Lin Niu, Zhihang Yuan, Dawei Yang, Xinggang Wang, and Wenyu Liu. Pd-quant: Post-training quantization based on prediction difference metric. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 24427–24437, 2023. 3, 6
- [33] Shunchang Liu, Jiakai Wang, Aishan Liu, Yingwei Li, Yijie Gao, Xianglong Liu, and Dacheng Tao. Harnessing perceptual adversarial patches for crowd counting. In *ACM CCS*, 2022. 1
- [34] Youquan Liu, Lingdong Kong, Jun Cen, Runnan Chen, Wenwei Zhang, Liang Pan, Kai Chen, and Ziwei Liu. Segment any point cloud sequences by distilling vision foundation models. *arXiv preprint arXiv:2306.09347*, 2023. 2
- [35] Zechun Liu, Baoyuan Wu, Wenhan Luo, Xin Yang, Wei Liu, and Kwang-Ting Cheng. Bi-real net: Enhancing the performance of 1-bit cnns with improved representational capability and advanced training algorithm. In *Proceedings of the European conference on computer vision (ECCV)*, pages 722–737, 2018. 8
- [36] Jun Ma and Bo Wang. Segment anything in medical images. *arXiv preprint arXiv:2304.12306*, 2023. 2
- [37] Yuexiao Ma, Huixia Li, Xiawu Zheng, Xuefeng Xiao, Rui Wang, Shilei Wen, Xin Pan, Fei Chao, and Rongrong Ji. Solving oscillation problem in post-training quantization through a theoretical perspective. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7950–7959, 2023. 3
- [38] Alaa Maalouf, Ninad Jadhav, Krishna Murthy Jatavallabhula, Makram Chahine, Daniel M Vogt, Robert J Wood, Antonio Torralba, and Daniela Rus. Follow anything: Open-set detection, tracking, and following in real-time. *arXiv preprint arXiv:2308.05737*, 2023. 1
- [39] Jeffrey L McKinstry, Steven K Esser, Rathinakumar Appuswamy, Deepika Bablani, John V Arthur, Izzet B Yildiz, and Dharmendra S Modha. Discovering low-precision networks close to full-precision networks for efficient inference. In *2019 Fifth Workshop on Energy Efficient Machine Learning and Cognitive Computing-NeurIPS Edition (EMC2-NIPS)*, pages 6–9. IEEE, 2019. 2
- [40] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1):99–106, 2021. 2
- [41] Daisuke Miyashita, Edward H Lee, and Boris Murmann. Convolutional neural networks using logarithmic data representation. *arXiv preprint arXiv:1603.01025*, 2016. 2, 8
- [42] Markus Nagel, Rana Ali Amjad, Mart Van Baalen, Christos Louizos, and Tijmen Blankevoort. Up or down? adaptive rounding for post-training quantization. In *International Conference on Machine Learning*, pages 7197–7206. PMLR, 2020. 1, 3, 6, 7
- [43] Markus Nagel, Marios Fournarakis, Rana Ali Amjad, Yelysei Bondarenko, Mart Van Baalen, and Tijmen Blankevoort. A white paper on neural network quantization. *arXiv preprint arXiv:2106.08295*, 2021. 1, 2
- [44] Lin Niu, Jiawei Liu, Zhihang Yuan, Dawei Yang, Xinggang Wang, and Wenyu Liu. Improving post-training quantization on object detection with task loss-guided lp metric. *CoRR*, 2023. 2
- [45] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *Advances in neural information processing systems*, 28, 2015. 6
- [46] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695, 2022. 2
- [47] David W Scott. Multivariate density estimation: Theory, practice and visualisation. john wiley and sons. Inc., New York, 1992. 4
- [48] Yuzhang Shang, Zhihang Yuan, Bin Xie, Bingzhe Wu, and Yan Yan. Post-training quantization on diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1972–1981, 2023. 3
- [49] Qihong Shen, Xingyi Yang, and Xinchao Wang. Anything-3d: Towards single-view anything reconstruction in the wild. *arXiv preprint arXiv:2304.10261*, 2023. 1, 2
- [50] Lv Tang, Haoke Xiao, and Bo Li. Can sam segment anything? when sam meets camouflaged object detection. *arXiv preprint arXiv:2304.04709*, 2023. 1
- [51] Zhi Tian, Chunhua Shen, Hao Chen, and Tong He. Fcos: Fully convolutional one-stage object detection. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9627–9636, 2019. 6
- [52] Jiakai Wang, Aishan Liu, Zixin Yin, Shunchang Liu, Shiyu Tang, and Xianglong Liu. Dual attention suppression attack: Generate adversarial camouflage in physical world. In *CVPR*, 2021. 1
- [53] Teng Wang, Jinrui Zhang, Junjie Fei, Yixiao Ge, Hao Zheng, Yunlong Tang, Zhe Li, Mingqi Gao, Shanshan Zhao, Ying Shan, et al. Caption anything: Interactive image description with diverse multimodal controls. *arXiv preprint arXiv:2305.02677*, 2023. 2
- [54] Yonghui Wang, Wengang Zhou, Yunyao Mao, and Houqiang Li. Detect any shadow: Segment anything for video shadow detection. *arXiv preprint arXiv:2305.16698*, 2023. 1
- [55] Ziwei Wang, Ziyi Wu, Jiwen Lu, and Jie Zhou. Bidet: An efficient binarized object detector. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2049–2058, 2020. 8

- [56] Xiuying Wei, Ruihao Gong, Yuhang Li, Xianglong Liu, and Fengwei Yu. Qdrop: Randomly dropping quantization for extremely low-bit post-training quantization. *arXiv preprint arXiv:2203.05740*, 2022. [1](#), [3](#), [6](#), [7](#)
- [57] Xiuying Wei, Yunchen Zhang, Xiangguo Zhang, Ruihao Gong, Shanghang Zhang, Qi Zhang, Fengwei Yu, and Xianglong Liu. Outlier suppression: Pushing the limit of low-bit transformer language models. *Advances in Neural Information Processing Systems*, 35:17402–17414, 2022. [1](#), [3](#)
- [58] Xiuying Wei, Yunchen Zhang, Yuhang Li, Xiangguo Zhang, Ruihao Gong, Jinyang Guo, and Xianglong Liu. Outlier suppression+: Accurate quantization of large language models by equivalent and optimal shifting and scaling. *arXiv preprint arXiv:2304.09145*, 2023. [1](#), [3](#)
- [59] Hao Wu, Patrick Judd, Xiaojie Zhang, Mikhail Isaev, and Paulius Micikevicius. Integer quantization for deep learning inference: Principles and empirical evaluation. *arXiv preprint arXiv:2004.09602*, 2020. [1](#), [6](#), [7](#)
- [60] Gui-Song Xia, Xiang Bai, Jian Ding, Zhen Zhu, Serge Belongie, Jiebo Luo, Mihai Datcu, Marcello Pelillo, and Liangpei Zhang. Dota: A large-scale dataset for object detection in aerial images. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3974–3983, 2018. [6](#)
- [61] Guangxuan Xiao, Ji Lin, Mickael Seznec, Hao Wu, Julien Demouth, and Song Han. Smoothquant: Accurate and efficient post-training quantization for large language models. In *International Conference on Machine Learning*, pages 38087–38099. PMLR, 2023. [1](#), [3](#)
- [62] Enze Xie, Wenhai Wang, Zhiding Yu, Anima Anandkumar, Jose M Alvarez, and Ping Luo. Segformer: Simple and efficient design for semantic segmentation with transformers. *Advances in Neural Information Processing Systems*, 34:12077–12090, 2021. [6](#)
- [63] Jinyu Yang, Mingqi Gao, Zhe Li, Shang Gao, Fangjing Wang, and Feng Zheng. Track anything: Segment anything meets videos. *arXiv preprint arXiv:2304.11968*, 2023. [1](#), [2](#)
- [64] Jingfeng Yao, Xinggang Wang, Lang Ye, and Wenyu Liu. Matte anything: Interactive natural image matting with segment anything models. *arXiv preprint arXiv:2306.04121*, 2023. [1](#)
- [65] Zhewei Yao, Reza Yazdani Aminabadi, Minjia Zhang, Xiaoxia Wu, Conglong Li, and Yuxiong He. Zeroquant: Efficient and affordable post-training quantization for large-scale transformers. *Advances in Neural Information Processing Systems*, 35:27168–27183, 2022. [3](#)
- [66] Tao Yu, Runseng Feng, Ruoyu Feng, Jinming Liu, Xin Jin, Wenjun Zeng, and Zhibo Chen. Inpaint anything: Segment anything meets image inpainting. *arXiv preprint arXiv:2304.06790*, 2023. [1](#), [2](#)
- [67] Zhihang Yuan, Chenhao Xue, Yiqi Chen, Qiang Wu, and Guangyu Sun. Ptq4vit: Post-training quantization for vision transformers with twin uniform quantization. In *European Conference on Computer Vision*, pages 191–207. Springer, 2022. [1](#), [2](#), [5](#), [6](#)
- [68] Chaoning Zhang, Dongshen Han, Yu Qiao, Jung Uk Kim, Sung-Ho Bae, Seungkyu Lee, and Choong Seon Hong. Faster segment anything: Towards lightweight sam for mobile applications. *arXiv preprint arXiv:2306.14289*, 2023. [2](#)
- [69] Hao Zhang, Feng Li, Shilong Liu, Lei Zhang, Hang Su, Jun Zhu, Lionel M Ni, and Heung-Yeung Shum. Dino: Detr with improved denoising anchor boxes for end-to-end object detection. *arXiv preprint arXiv:2203.03605*, 2022. [6](#), [7](#)
- [70] Renrui Zhang, Zhengkai Jiang, Ziyu Guo, Shilin Yan, Junting Pan, Hao Dong, Peng Gao, and Hongsheng Li. Personalize segment anything model with one shot. *arXiv preprint arXiv:2305.03048*, 2023. [1](#)
- [71] Xu Zhao, Wenchao Ding, Yongqi An, Yinglong Du, Tao Yu, Min Li, Ming Tang, and Jinqiao Wang. Fast segment anything. *arXiv preprint arXiv:2306.12156*, 2023. [2](#), [8](#)
- [72] DanDan Zheng, Yuanliu Liu, Liang Li, et al. Leveraging inter-layer dependency for post-training quantization. *Advances in Neural Information Processing Systems*, 35:6666–6679, 2022. [2](#)
- [73] Bolei Zhou, Hang Zhao, Xavier Puig, Sanja Fidler, Adela Barriuso, and Antonio Torralba. Scene parsing through ade20k dataset. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 633–641, 2017. [6](#)
- [74] Xueyan Zou, Jianwei Yang, Hao Zhang, Feng Li, Linjie Li, Jianfeng Gao, and Yong Jae Lee. Segment everything everywhere all at once. *arXiv preprint arXiv:2304.06718*, 2023. [2](#)