

# Distribution Adaptive INT8 Quantization for Training CNNs

Kang Zhao, Sida Huang, Pan Pan, Yinghan Li, Yingya Zhang, Zhenyu Gu, Yinghui Xu

Machine Intelligence Technology Lab, Alibaba Group

{zhaokang.zk, sida.hsd, panpan.pp, lyh238099, yingya.zyy, zhenyu.gu, renji.xyh}@alibaba-inc.com

## Abstract

Researches have demonstrated that low bit-width (e.g., INT8) quantization can be employed to accelerate the inference process. It makes the gradient quantization very promising since the backward propagation requires approximately twice more computation than forward one. Due to the variability and uncertainty of gradient distribution, a lot of methods have been proposed to attain training stability. However, most of them ignore the channel-wise gradient distributions and the impact of gradients with different magnitudes, resulting in the degradation of final accuracy. In this paper, we propose a novel INT8 quantization training framework for convolutional neural network to address the above issues. Specifically, we adopt *Gradient Vectorized Quantization* to quantize the gradient, based on the observation that layer-wise gradients contain multiple distributions along the channel dimension. Then, *Magnitude-aware Clipping Strategy* is introduced by taking the magnitudes of gradients into consideration when minimizing the quantization error, and we present a theoretical derivation to solve the quantization parameters of different distributions. Experimental results on broad range of computer vision tasks, such as image classification, object detection and video classification, demonstrate that the proposed Distribution Adaptive INT8 Quantization training method has achieved almost lossless training accuracy for different backbones, including ResNet, MobileNetV2, InceptionV3, VGG and AlexNet, which is superior to the state-of-the-art techniques. Moreover, we further implement the INT8 kernel that can accelerate the training iteration more than 200% under the latest Turing architecture, i.e., our method excels on both training accuracy and speed.

## Introduction

Many literatures (Jacob et al. 2018; Wang et al. 2018b; Zhuang et al. 2018; Liu et al. 2018; Jain et al. 2019) have demonstrated that low bit-width quantization can be utilized to speed up the inference stage (weight or activation) with minor hurt of accuracy. It makes quantization of gradient quite promising since the backward propagation consumes twice more computing resources than forward propagation (Banner et al. 2018). Among the family of low bit-width formats (e.g., FP16, INT16, FP8, INT8, INT4, etc), INT8 is widely adopted in community for the trade-off between

Copyright © 2021, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

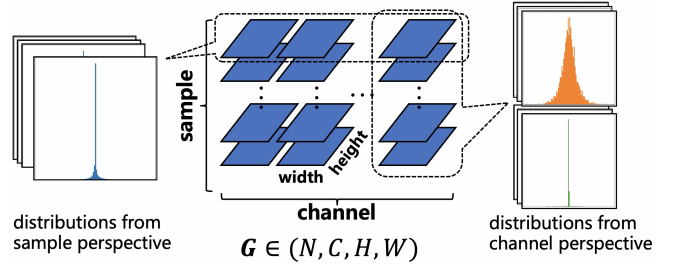


Figure 1: The observation of layer-wise gradient distribution from sample and channel perspectives, with respect to conv6, 21, 34, 46 of ResNet-50 on ImageNet. It is obvious there exist two categories of gradient distribution observed along channel dimension: a bell-shaped distribution, and a sharp with long-tailed shape distribution. But the distributions observed from sample perspective look similar.

efficiency and accuracy. On the one hand, INT8 operation is twice faster than FP16/INT16, and supported by more architectures compared with FP8. On the other hand, INT8 has a larger quantization range (-128~127) than INT4 (-8~7) or other lower bit-widths (less than 4-bits).

Although quantizing gradient with INT8 to accelerate training is appealing, there still remains a big challenge, because the quantization error of gradient easily misleads the direction of convergence as analyzed in (Zhu et al. 2020). Many methods have been proposed to keep the training process reliable, most of them either use floating point numbers to simulate INT8 calculation, which has no actual acceleration effect, or resort to inappropriate quantization schemes that result in large deterioration of final accuracy. (Zhu et al. 2020) adopts a unified solution to make INT8 training stable for a variety of networks, but it still suffers from some loss of accuracy (up to 1% on ImageNet), which is not acceptable in practice. Combining INT16 with INT8 (Zhang et al. 2020) is another option to compensate the quantization loss of gradient but outside the scope of our work.

In this paper, we propose a Distribution Adaptive INT8 Quantization training framework from the aspects of training accuracy and speedup. The main contributions of our work are outlined as follows:

- Based on an important observation ignored by previous

works that the layer-wise gradients contains multiple distributions, we employ *Gradient Vectorized Quantization* to well capture the gradient distribution.

- We propose *Magnitude-aware Clipping Strategy* which considers the impact of gradients with different magnitudes when minimizing quantization error, and present a theoretical derivation to solve the optimal quantization parameters of different distributions.
- Under the latest Turing architecture, we implement TensorCore<sup>1</sup>-based INT8 kernel, which achieves 18% faster than FP16 and speeds up FP32 more than 200%.
- To the best of our knowledge, we are the first INT8 training framework to achieve almost lossless performance compared with full precision (FP32) on a variety of tasks and backbones.

## Related Work

Generally speaking, quantization methods can be roughly divided into two major categories: inference quantization and training quantization, according to distinct quantization targets.

**Inference Quantization.** Inference quantization aims to quantize weight and/or activation to accelerate the forward pass (Nagel et al. 2019; Zhang et al. 2019). BNNs (Hubara et al. 2016) trains the neural networks with binary weight and activation successfully. TWNs (Li, Zhang, and Liu 2016) constrains weight to +1, 0 and -1 and optimizes a threshold-based ternary function. These low bit-widths (less than 4-bits) quantization techniques have significant degradation in performance. To improve the accuracy, more advanced quantization schemes (Cai et al. 2017; Zhang et al. 2018; Jung et al. 2019) have been proposed. In (Choi et al. 2019), the combination of range ReLU and statistics-aware weight binning make the 2-bit quantization get comparable performance with full precision networks. By training quantization thresholds, TQT (Jain et al. 2019) also achieves near-floating-point accuracy.

**Training Quantization.** Compared with inference quantization, training quantization needs to further quantize the gradient to accelerate the whole training process. (Banner et al. 2018) replaces batch-norm (BN) layer with RangeBN and adopt “Gradients Bifurcation” to quantize both gradient and BN layers to 8-bit, but it de-quantizes gradient to floating point number before calculation, which can not be speed up by INT8 operator. A similar case is (Wang et al. 2018a), which needs new hardware platforms to boost the speed. Besides quantizing gradient to INT8, WAGE (Wu et al. 2018) and WAGEUBN (Yang et al. 2020) further quantize other training modules (e.g., Error and Update), leading to large loss of accuracy. (Zhu et al. 2020) utilizes direction sensitive gradient clipping and counteractive learning rate scaling to guide the 8-bit gradient learning and updating, which help it succeed in many networks and tasks.

Considering the INT8 quantization of weight and activation have been well studied in (Zhu et al. 2020; Zhang et al. 2020), we will take the same way to quantize weight and

activation as they did, and concentrate on the gradient quantization in the rest of the paper.

## Our Method

This section describes the formulation of our INT8 quantization framework. To facilitate our discussion, the quantization method is given below.

Following (Zhu et al. 2020; Zhang et al. 2020), we employ the symmetric uniform quantization. For a given data  $x$ , and a quantization parameter  $s$ , the fixed point 8-bit integer quantization is defined as:

$$q(x) = \text{round}(127 \cdot \frac{\text{clamp}(x, s)}{s}), \quad (1)$$

where  $\text{clamp}(\cdot)$  is a clipping function:

$$\text{clamp}(x, s) = \begin{cases} x, & |x| \leq s \\ \text{sign}(x) \cdot s, & |x| > s \end{cases} \quad (2)$$

It can guarantee  $x$  in the range of  $[-s, s]$ . Correspondingly, the de-quantization of  $x$  can be easily expressed as:  $\hat{x} = q(x) \times \frac{s}{127}$ . Similar to prevailing methods, we utilize stochastic rounding (Gupta et al. 2015) for gradient quantization and nearest rounding for weight and activation quantization.

## Motivation

Clearly,  $s$  is critical to the quantization error of gradient, since it is the only variable in Eq.(1). The existing approaches either simply take  $s$  as  $\max(|x|)$  or minimize the mean-square error (MSE) or cosine distance between  $x$  and  $\hat{x}$  to find a better value of  $s$ . They all encounter the following two issues:

i) **The multiple distributions of the layer-wise gradients are not taken into account.** In order to minimize the quantization error, the previous ways (Banner et al. 2018; He and Fan 2019; Zhu et al. 2020; Zhang et al. 2020) try to find a global value of  $s$  by assuming that the gradient satisfies a certain distribution, or adopt periodic gradient descent trick for seeking an optimal  $s$  without distribution assumption. These methods only use one global quantization parameter ( $s$ ) for each layer’s gradients (we call it global quantization for short). Based on our observations, we found that there exist more than one distribution of gradients inside one layer. And different distributions usually have different optimal quantization parameters. For example, taking  $s$  as  $\max(|x|)$  is good enough for Gaussian Distribution, yet Laplacian Distribution will prefer a  $s$  smaller than  $\max(|x|)$  in consideration of the densely distributed gradients that are far away from the maximum. It indicates multiple quantization parameters can capture the gradient distribution better than only one quantization parameter.

ii) **The contribution of gradients with different magnitudes is not considered.** (Lin et al. 2017b; Strom 2015) show that it is more important to transmit large gradients than small ones during the communication process, because large gradients contain more information. Similarly, the quantization error of large gradients should be more important than small ones in view of training accuracy. As stated in (Banner et al. 2018; Zhu et al. 2020), small values

<sup>1</sup><https://www.nvidia.com/en-us/data-center/tensorcore/>

will take the majority of gradients with training evolving, which means if we ignore the magnitudes of gradients, the quantization error and quantization parameter will be dominated by the small gradients. In that case, the quantization error of large gradients will be larger, leading to the deterioration of final accuracy.

To solve these two issues, we propose *Gradient Vectorized Quantization* and *Magnitude-aware Clipping Strategy*, respectively.

### Gradient Vectorized Quantization

Without loss of generality, we assume the input of a convolution layer is  $\mathbf{X} \in (N, C_{in}, H_{in}, W_{in})^2$ .  $\mathbf{W} \in (C_{out}, C_{in}, k_1, k_2)$  is the weight parameter and the output is  $\mathbf{Y} \in (N, C_{out}, H_{out}, W_{out})$ . During backward propagation, the gradients of  $\mathbf{Y}$  is denoted as  $\mathbf{G}$ . We observe  $\mathbf{G}$  of each layer from both sample and channel perspectives during the training process, as shown in Figure 1. It is found that the channel dimension obviously split the distribution of  $\mathbf{G}$  into two categories: a bell-shaped distribution, and a sharp with long-tailed shape distribution (we call it Inverted-T distribution). Similar phenomena is not found from the sample dimension, which may arise from the specific attributes contained in channel dimension. Considering we will use different quantization parameters for these two distributions (details shown in the next section), we apply *Gradient Vectorized Quantization* along the channel dimension.

In the gradient calculation of convolution layer,  $\mathbf{G}$  is coming from the previous layer, then  $\mathbf{G}_X$  and  $\mathbf{G}_W$  can be calculated as following:

$$\mathbf{G}_X = \mathbf{G} \odot \mathbf{W}, \quad (3)$$

$$\mathbf{G}_W = \mathbf{X}' \otimes \mathbf{G}', \quad (4)$$

where  $\odot$  and  $\otimes$  are two specific convolution operations<sup>3</sup>. In the implementation of Eq.(4), we usually do a dimension transformation on  $\mathbf{X}$  and  $\mathbf{G}$  to make the shape match:  $\mathbf{X} \rightarrow \mathbf{X}' \in (C_{in}, N, H_{in}, W_{in})$ ,  $\mathbf{G} \rightarrow \mathbf{G}' \in (C_{out}, N, H_{out}, W_{out})$ .

Now, we consider the vectorized quantization of Eq.(4).  $q(\mathbf{X}')$  can be easily obtained by making a similar dimension transformation of  $q(\mathbf{X})$  (it has been calculated in forward pass), and its quantization parameter is denoted as  $s_x$ . Let  $\mathbf{G}' = [\mathbf{G}'_1, \dots, \mathbf{G}'_i, \dots, \mathbf{G}'_{C_{out}}]^\top$ , where  $\mathbf{G}'_i \in (N, H_{out}, W_{out})$ , we quantize  $\{\mathbf{G}'_i\}_{i=1}^{C_{out}}$  with different quantization parameter  $\{s_i\}_{i=1}^{C_{out}}$  (how to get  $s_i$  will be described in the following section):

$$q(\mathbf{G}') = [q(\mathbf{G}'_1), \dots, q(\mathbf{G}'_i), \dots, q(\mathbf{G}'_{C_{out}})]^\top. \quad (5)$$

Then we have:

$$\begin{aligned} q(\mathbf{G}_W) &= q(\mathbf{X}') \otimes q(\mathbf{G}') \\ &= [q(\mathbf{X}') \otimes q(\mathbf{G}'_1), \dots, q(\mathbf{X}') \otimes q(\mathbf{G}'_i), \dots]^\top \\ &= [q(\mathbf{G}_{W_1}), \dots, q(\mathbf{G}_{W_i}), \dots, q(\mathbf{G}_{W_{C_{out}}})]^\top, \end{aligned}$$

<sup>2</sup> $(N, C_{in}, H_{in}, W_{in})$  is short for  $\mathbb{R}^{N \times C_{in} \times H_{in} \times W_{in}}$

<sup>3</sup> $\odot$  is transpose convolution, and  $\otimes$  is dilation convolution in practice

where  $q(\mathbf{G}_{W_i}) = q(\mathbf{X}') \otimes q(\mathbf{G}'_i)$  with the shape size  $(C_{in}, k_1, k_2)$ . The de-quantization of  $\mathbf{G}_W$  is also very simple:

$$\hat{\mathbf{G}}_W = [\hat{\mathbf{G}}_{W_1}, \dots, \hat{\mathbf{G}}_{W_i}, \dots, \hat{\mathbf{G}}_{W_{C_{out}}}]^\top, \quad (6)$$

$$\hat{\mathbf{G}}_{W_i} = q(\mathbf{G}_{W_i}) \cdot \frac{s_x}{127} \cdot \frac{s_i}{127}, \quad (7)$$

which can be implemented by an elementary product operation in PyTorch framework.

We use global quantization (as we did in the forward pass) to quantize Eq.(3), since there is no obvious difference in the distribution of  $\mathbf{G}$  from sample dimension, which is also consistent with (Banner et al. 2018) that  $\mathbf{G}_W$  need higher quantization accuracy than  $\mathbf{G}_X$ .

Noted that although vectorized quantization has  $C$  (numbers of channel) quantization parameters, it has the same computational complexity as global quantization. Take seeking the quantization parameter as an example, global quantization requires searching  $N \times C \times H \times W$  float spaces to find the optimal  $s$ , while vectorized quantization only needs to search  $N \times H \times W$  float spaces for each quantized parameter, which is basically equal on the whole.

### Magnitude-aware Clipping Strategy

Because large gradients contribute more on network accuracy than smaller ones, we introduce the magnitudes of gradients ( $f(g)$ ) into quantization error as follows:

$$E = \int_{g_{min}}^{g_{max}} |g - \hat{g}| f(g) p(g) dg, \quad (8)$$

where  $g$  is the continuous random variable in the range of  $[g_{min}, g_{max}]$  for  $\{\mathbf{G}'_i\}_{i=1}^{C_{out}}$ ,  $\hat{g}$  stands for its corresponding de-quantization, and  $p(g)$  denotes the gradient distribution.

Mathematically, the positive correlation between magnitude contribution and gradient amplitude ( $|g|$ ) can be expressed as  $f(g) = e^{\alpha|g|}$ . Note that nonnegative  $\alpha$  reflects the importance of large amplitude to quantization error. More concretely, for two gradients  $g_1, g_2$  with  $|g_1| < |g_2|$ ,  $frac(E, g_2)$ <sup>4</sup> becomes much greater than  $frac(E, g_1)$  with the increase of  $\alpha$ .

Considering the complexity of Eq.(8) and difficult-to-measure of  $\alpha$ , we make the following two assumptions to grasp the key features in choosing quantization parameter:

**Assumption 1**  $\forall g \in \mathbb{R}, p(g) = p(-g)$ ;

**Assumption 2**  $\forall g \in (0, s), |g - \hat{g}| \approx \frac{1}{2} \cdot \frac{s}{127}$ .

The first assumption is based on our observation that the gradient distribution is nearly symmetric around the origin. And the second one utilizes the rectangle rule, which is widely adopted in numerical integration (McCallum et al. 2012). More empirical results are presented in Appendix to support these assumptions. Then, we can simplify Eq.(8) to:

$$E = \underbrace{\int_0^s \frac{s}{127} \cdot f(g) p(g) dg}_{I_1} + 2 \underbrace{\int_s^{g_{max}} (g - s) \cdot f(g) p(g) dg}_{I_2}, \quad (9)$$

<sup>4</sup> $frac(E, x) = \frac{1}{E} \int_x^{x+\epsilon} |g - \hat{g}| f(g) p(g) dg$ , where  $\epsilon$  is an infinitesimal value

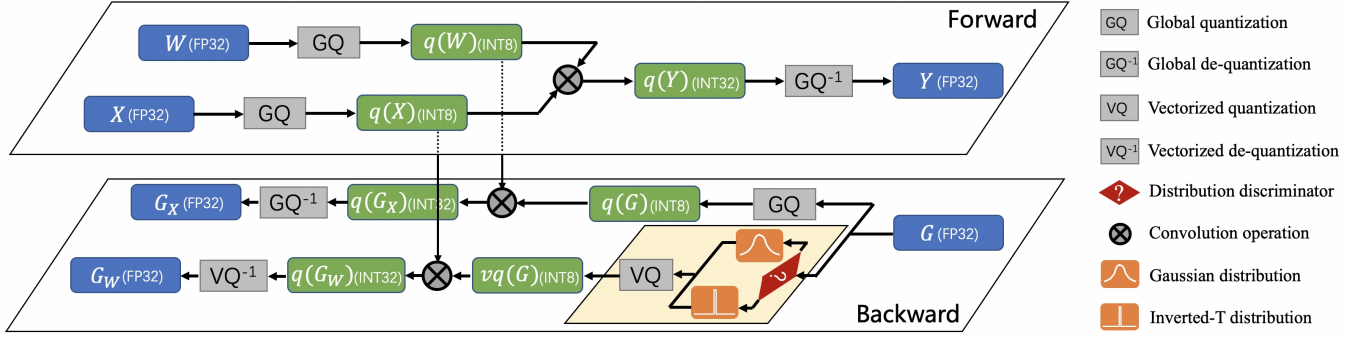


Figure 2: The dataflow of Distribution Adaptive INT8 Quantization for convolution layer. In the forward pass, we use global quantization to quantize weight and activation. While in the backward process, we adopt vectorized quantization of gradient to calculate  $G_W$ , and global quantization for  $G_X$ .

where  $s \leq g_{max}$  is obvious. Eq.(9) quantitatively describes the contradictory effect of  $s$  on  $E$ : as  $s$  goes down,  $I1$  decreases yet  $I2$  increases, and vice versa. To seek a balance between  $I1$  and  $I2$ , the optimal quantization parameter  $\bar{s}$  should satisfy the following condition:

$$\left. \frac{\partial E}{\partial s} \right|_{s=\bar{s}} = 0. \quad (10)$$

Next, we discuss how to choose  $\bar{s}$  under different gradient distributions.

**Gaussian distribution.** We might as well assume the bell-shaped gradients can be described by Gaussian distribution, which takes the max absolute value ( $|g|_{max}$ ) as optimal quantization parameter (Banner et al. 2018):

$$\bar{s} = |g|_{max}. \quad (11)$$

In general, the quantization parameter can be denoted by:  $s = \beta |g|_{max}$ , with  $\beta \in (0, 1]$ . There exists the following inequality constraint between  $\alpha$  and  $\beta$ :

$$\frac{\partial}{\partial \alpha} \left( \frac{\partial E}{\partial s} \right) \cdot \frac{\partial}{\partial \beta} \left( \frac{\partial E}{\partial s} \right) < 0. \quad (12)$$

Eq.(12) suggests that  $\beta$  increases as  $\alpha$  grows. The increase of  $\beta$  makes  $s$  closer to  $|g|_{max}$ , so we reuse Eq.(11) to get quantization parameter of Gaussian distribution after considering the magnitude contribution ( $\alpha$ ).

**Inverted-T distribution.** As shown in the Figure 1, the majority of gradients are concentrated near zero, while those far away from the center have much smaller density. This extremely unbalanced distribution can be well portrayed by a piecewise uniform distribution, namely:

$$p_{T^{-1}}(g) = \begin{cases} a, & |g| \in (0, \epsilon) \\ b, & |g| \in (\epsilon, |g|_{max}) \end{cases}, \quad (13)$$

where  $a$  and  $b$  determines distribution intensity with  $a \gg b$ , and  $\epsilon$  is a small value. Note that  $|p(g) - p_{T^{-1}}(g)|$  is trivial comparing to the large value of  $a$ .

Combining Eq.(13) and Eq.(9), the first order derivative of  $E$  with respect to  $s$  is reduced to:  $\frac{\partial E}{\partial s} =$

$\frac{(a-b)e^{\alpha\epsilon} + b(255+\alpha s)e^{\alpha s} - 254be^{\alpha|g|_{max}} - a}{127\alpha}$ . To eliminate the influence of  $\alpha$ , we make Eq.(10) of adjacent iterations equal:  $\frac{\partial E}{\partial s} \Big|_{s=\bar{s}_t} = \frac{\partial E}{\partial s} \Big|_{s=\bar{s}_{t-1}} = 0$ . Then we can derive the following iterative formula:

$$\bar{s}_t = (1 - kA)\bar{s}_{t-1} + A|g|_{max,t}, \quad (14)$$

where  $k$  and  $A$  are the hyper-parameters (The proof of Eq.(12) and Eq.(14) are presented in Appendix). Eq.(14) indicates that the gradients calculated from a mini-batch can be improved by the complete dataset after sufficient iterations. Therefore, the update formula in Eq.(14) may have sort of regularization effect on the gradient.

### Distribution Adaptive INT8 Quantization

Finally, we show how to distinguish different distributions along channel dimension. Based on our observation, Inverted-T distribution ( $T^{-1}$ ) has more values concentrated around zero than Gaussian distribution. Thus, we propose the following gradient discriminator:

$$\begin{cases} p \sim N(\mu, \sigma^2), P(|g| > \sigma) > \lambda \\ p \sim T^{-1}, P(|g| > \sigma) \leq \lambda \end{cases}, \quad (15)$$

where  $\mu$  is the mean of gradient. The threshold  $\lambda$  is set as 0.3 in this work, because more than 30% of values drawn from a Gaussian distribution are outside  $\sigma$  (Patel and Read 1996) theoretically.

We summarize the complete dataflow of applying our INT8 quantization for convolution layer in Figure 2. Two quantizations of gradient (vectorized + global) can be easily optimized with kernel fusion in implementation (see Appendix for more details). To highlight the core pipeline of quantizing gradient, we describe the backward pass of Distribution Adaptive INT8 Quantization in Algorithm 1.

### Experiments

In this section, we evaluate the proposed method on different tasks and test the training speedup in practice. We use cosine scheduler (Loshchilov and Hutter 2017) for MobileNetV2 (Sandler et al. 2018), and multistep learning rate scheduler for the others. All experimental settings for INT8 training

**Algorithm 1** Backward pass of Distribution Adaptive INT8 Quantization.

**Input:** Full precision gradient of  $l$ -layer  $G_{Y_l}$  in  $t$  training iteration with shape  $(N, C, H, W)$ , quantized activation and weight saved in forward pass  $q(X_{l-1})$  and  $q(W_l)$ , and their quantization parameters  $s_{X_{l-1}}$  and  $s_{W_l}$ .

**Output:**  $G_{X_{l-1}}$  and  $G_{W_l}$  // gradients for  $X_{l-1}$  and  $W_l$

```

1: for  $c \leftarrow 1$  to  $C$  do
2:   Infer  $|g|_{max,l,c}$  and  $\sigma_{l,c}$ ;
3:   // density of gradients larger than  $\sigma_{l,c}$ 
4:    $P = \sum_{|g| > \sigma_{l,c}} \frac{1}{N \cdot H \cdot W}$ ;
5:   if  $P > \lambda$  then // update as Gaussian distribution
6:      $s_{l,c,t} = |g|_{max,l,c}$ ;
7:   else // update as Inverted-T distribution
8:      $s_{l,c,t} = (1 - kA)s_{l,c,t-1} + A|g|_{max,l,c}$ ;
9:   end if
10: end for
11:  $vq(G_{Y_l}) = VQ(G_{Y_l}, s_l)$ ; // vectorized quantization
12:  $q(G_{Y_l}) = Q(G_{Y_l}, |g|_{max,l})$ ; // global quantization
13:  $q(G_{W_l}) = q(X'_{l-1}) \otimes vq(G'_{Y_l})$ ; // dilation convolution
14:  $q(G_{X_{l-1}}) = q(G_{Y_l}) \odot q(W_l)$ ; // transpose convolution
15: // apply vectorized/global de-quantization
16:  $G_{W_l} = VQ^{-1}(q(G_{W_l}), s_{W_l}, s_l)$ ;
17:  $G_{X_{l-1}} = Q^{-1}(q(G_{X_{l-1}}), s_{X_{l-1}}, |g|_{max,l})$ .

```

are consistent with the full precision model (baseline). Due to the mismatched baselines between different works, we adopt the accuracy drop ( $\Delta = \text{INT8} - \text{baseline}$ ) as a metric if not specified.

### Ablation study

The ablation study is conducted on ImageNet (Deng et al. 2009) with ResNet-18 (He et al. 2016). We set the global quantization (GQ) of gradient as baseline, which takes  $|g|_{max}$  as quantization parameter.

**Gradient Vectorized Quantization (GVQ).** To verify the effectiveness of vectorized quantization, we use the identical strategy of GQ to update the quantization parameter along channel dimension (denoted as Ours<sub>GVQ</sub>). Figure 3 (a) presents the evolution of quantization errors  $E$  in different epochs. We can find the quantization error of Ours<sub>GVQ</sub> is significantly lower than that of GQ during the total training process, which shows the quantization strategy of treating each channel differently is more reasonable than GQ.

**Magnitude-aware Clipping Strategy (MCS).** In the second experiment, we apply MCS to the layer-wise (not channel-wise) gradients based on Eq.(15) to get free of GVQ (denoted as Ours<sub>MCS</sub>). We utilize Kolmogorov-Smirnov (KS) statistic (Kolmogorov 1933) to quantify the difference between two distributions:

$$D_n = \sup_g |F_1(g) - F_2(g)|, \quad (16)$$

where  $\sup_g$  is the supremum of the set of distances,  $F_1$  is the cumulative distribution function of gradient, and  $F_2$  is

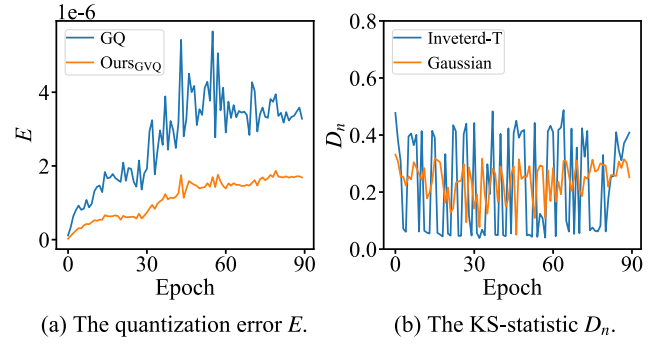


Figure 3: The evolution curves of quantization error (with  $\alpha = 0.2$  for representative) and KS-statistic for a convolution layer in ResNet-18.

$k$	1.0	1.0	1.2	1.2	1.5	1.5
$A$	0.5	0.8	0.5	0.8	0.5	0.8
Acc./%	69.94	70.01	69.90	69.82	69.80	69.91

Table 1: Comparison of different  $k$  and  $A$ .

for the approximation one (Gaussian/Inverted-T). Figure 3 (b) shows the  $D_n$  of Inverted-T (blue curve) is lower than that of Gaussian (orange curve) in some intervals of training epochs, indicating the rationality of Inverted-T assumption. Compared with previous methods that assume Gaussian distribution of gradient, our distribution adaptive strategy can describe the ever-changing gradient more precisely.

Furthermore, we explore the sensitivity of two additional hyper-parameters ( $k$  and  $A$ ) introduced in MCS. As Table 1 lists, the accuracy between different hyper-parameters are similar, which proves the stability of our MCS. In the following experiments, we set  $k = 1$  and  $A = 0.8$ .

**Results.** Table 2 summarizes the final accuracy of different quantization schemes. Ours<sub>GVQ</sub> and Ours<sub>MCS</sub> obtain 0.30% and 0.33% accuracy improvement over GQ respectively, showing again these two strategies are effective. Not surprisingly, our proposed framework has the highest accuracy improvement, i.e., 0.53%, which exceeds both Ours<sub>GVQ</sub> and Ours<sub>MCS</sub>. This is attributed to the fact that our classification of gradient distribution from channel perspective has a smaller  $D_n$  value than that of layer perspective.

### Image Classification

Next, we present the classification performance (Song et al. 2020) on CIFAR-10 (Krizhevsky et al. 2009) and ImageNet datasets, with different backbones: AlexNet (Krizhevsky, Sutskever, and Hinton 2012), ResNet, VGG (Simonyan and

Method	GQ	Ours <sub>GVQ</sub>	Ours <sub>MCS</sub>	Ours
Acc / %	69.68	69.98	70.01	<b>70.21</b>

Table 2: Ablation study on different quantization strategies.

Model	Method	FP32/%	INT8/%	$\Delta$ /%
ResNet-20	UI8	92.32	91.95	-0.37
	Ours	92.35	92.76	<b>0.41</b>
MobileNetV2	UI8	94.39	93.38	-1.01
	Ours	94.73	94.37	<b>-0.36</b>
InceptionV3	UI8	94.89	95.00	0.11
	Ours	95.00	95.21	<b>0.21</b>

Table 3: The Top-1 Accuracy on CIFAR-10 dataset.

Zisserman 2015), MobileNetV2, and InceptionV3 (Szegedy et al. 2016). We compare our method with the following state-of-the-art: Unified Int8 training (UI8) (Zhu et al. 2020), Adaptive Fix-Point training (AFP) (Zhang et al. 2020), WAGEUBN (Yang et al. 2020), FP8 training (Wang et al. 2018a) and DoReFa-Net (Zhou et al. 2016).

**CIFAR-10.** It’s clear the proposed method has less accuracy drop than UI8 for every examined model, as shown in Table 3. For MobileNetV2, we even make the accuracy loss less than 0.4%, proving the effectiveness of our INT8 training framework. Our method achieves higher training accuracy than baseline for InceptionV3 just like UI8, but we extend it to ResNet-20 further.

**ImageNet.** Table 4 shows the top-1 accuracy of different backbones on ImageNet dataset. The performance of our INT8 quantization is superior to other state-of-the-art methods on all models. The quantization strategy of WAGEUBN is too aggressive (quantizing all data paths), resulting in considerable accuracy loss on many models (larger than 5% on ResNet-50). For AlexNet, DoReFa-Net has the worst performance due to its inappropriate quantization of gradient. FP8 training performs better than other methods on ResNet-18, but it depends on new hardware to boost the speed. For the popular ResNet-50, our method achieves even 0.09% accuracy improvement over baseline, while UI8 and AFP cause more than 0.2% accuracy drop, which verifies our multiple quantization parameters are more reasonable than only one. Although AFP has better performance in AlexNet than ours, their computation is more intensive owing to 77.5% INT16 calculation.

Because of the sparsely linked depth-wise convolution structure, the quantization of MobileNetV2 in previous works always cause a significant accuracy drop, i.e., 1.19% for UI8 and 1.30% for AFP. Equipped with our method, the accuracy drop immediately decreases to 0.52%, which is almost twice better than UI8 and AFP. Therefore, our consideration of the inherent gradient distributions and magnitudes can benefit the quantization of complex convolution layer.

To the best of our knowledge, except for MobileNetV2, we are the first INT8 framework to obtain almost lossless performance compared with full precision, which is hardly achieved by previous studies. It is very important in practical scenarios, because one prefers spending more time training the FP32 model if there are some loss of accuracy in the INT8 quantization scheme. Moreover, our performance on

Model	Method	FP32/%	INT8/%	$\Delta$ /%
Alexnet	UI8	59.84	58.84	-1.00
	AFP	58.00	58.22	<b>0.22</b>
	DoReFa-Net	55.90	53.00	-2.90
	FP8 training	58.00	57.50	-0.50
	Ours	58.18	58.21	0.03
ResNet-18	UI8	70.30	69.67	-0.63
	WAGEUBN	68.70	67.40	-1.30
	FP8 training	67.43	67.34	-0.09
	Ours	70.22	70.21	<b>-0.01</b>
ResNet-34	UI8	73.68	73.29	-0.39
	WAGEUBN	71.99	68.50	-3.49
	Ours	73.46	73.40	<b>-0.06</b>
ResNet-50	UI8	76.60	76.34	-0.26
	AFP	76.40	76.20	-0.20
	WAGEUBN	74.66	69.07	-5.59
	Ours	76.50	76.59	<b>0.09</b>
InceptionV3	UI8	72.39	71.20	-1.19
	AFP	73.00	72.80	-0.20
	Ours	75.47	75.48	<b>0.01</b>
VGG-16	AFP	71.00	70.60	-0.40
	Ours	72.39	72.44	<b>0.05</b>
MobileNetV2	UI8	72.39	71.20	-1.19
	AFP	71.80	70.50	-1.30
	Ours	72.44	71.92	<b>-0.52</b>

Table 4: The Top-1 Accuracy on ImageNet dataset.

AlexNet, VGG-16 and InceptionV3 are slightly better than the baseline, which may be attributed to the regularization effect originated from Eq.(14).

## Object Detection

To further validate the extensibility of our method, we train Faster-RCNN (Ren et al. 2015) and RetinaNet (Lin et al. 2017a) on COCO (Lin et al. 2014) and PASCAL VOC datasets (Everingham et al. 2010). We adopt the open-source MMDetection framework (Chen et al. 2019) with the default settings and take ResNet-50 as backbone. Mean Average Precision (mAP) is used as evaluation metric here.

As listed in Table 5, our method achieves satisfactory results on both datasets, which are less than 0.1% accuracy drop compared to full precision model. In particular, we surpasses UI8 on the large-scale COCO dataset by a large margin (more than 1%). Our good performance on both Faster-RCNN (two-stage) and RetinaNet (one-stage) also reveal the versatility of our method for different types of detectors. In contrast, UI8 suffers from dramatic accuracy drop (larger than 1%) on these two models. Note that we are the first to report almost lossless INT8 results on the object detection task with various models.



Dataset	Model	Method	FP32/%	INT8/%	$\Delta$ /%
COCO	Faster-RCNN	UI8	36.9	35.1	-1.8
		Ours	37.5	37.4	<b>-0.1</b>
	Retina-Net	UI8	36.2	35.0	-1.2
		Ours	36.4	36.3	<b>-0.1</b>
PASCAL-VOC	Faster-RCNN	UI8	82.0	81.9	-0.1
		Ours	79.5	79.4	<b>-0.1</b>
	Retina-Net	UI8	-	-	-
		Ours	77.3	77.4	<b>0.1</b>

Table 5: Detection Results on different datasets.

Dataset	Model	FP32/%	Ours/%	$\Delta$ /%
UCF-101	ResNet3D-18	54.9	55.1	<b>0.2</b>
	ResNet3D-50	57.6	57.5	<b>-0.1</b>
Kinetics-400	ResNet3D-18	51.0	51.4	<b>0.4</b>
	ResNet3D-50	58.5	58.8	<b>0.3</b>

Table 6: The Top-1 Accuracy of Video Classification.

## Video Classification

Applying quantization to 3D convolution is very meaningful owing to the high computation complexity of 3D convolution layer. The transition of our method from 2D to 3D convolution layer is straight forward: treating depth dimension in the same way as height and width. In addition, we found the optimal hyper-parameters ( $k$ ,  $A$ ) used in 2D layer performs well in 3D case, showing the stability of our quantization scheme again. We conduct experiments with ResNet3D-18/50 for video classification. Two widely used dataset, i.e., UCF-101 (Soomro, Zamir, and Shah 2012) with official split-1 and Kinetics-400 (Kay et al. 2017) released in 2017 are both tested.

Our INT8 quantization in all examined models perform comparable with (even better than) baseline in Table 6. The success of our method in 3D convolution demonstrates the similarity between depth and height/width components, which also reflects the importance of our analysis of gradient from the channel dimension.

## Acceleration Results

To make full use of the computational capacity of INT8, we take the advantage of TensorCore to implement the convolution kernel. We specially design three-stage pipelines: GPU memory  $\Rightarrow$  Shared memory  $\Rightarrow$  Register  $\Rightarrow$  TensorCore, to make computation and memory access fully parallel (Appendix contains more details).

In Figure 4, we compare our INT8 implementation over FP16 version (based on cuDNN 7.6.5) on various convolution layers of ResNet-50. The implementation algorithm is limited to IMPLICIT\_PRECOMP\_GEMM<sup>5</sup> for fair com-

<sup>5</sup>[https://docs.nvidia.com/deeplearning/cudnn/api/index.html#cudnnConvolutionFwdAlgo\\_t](https://docs.nvidia.com/deeplearning/cudnn/api/index.html#cudnnConvolutionFwdAlgo_t)

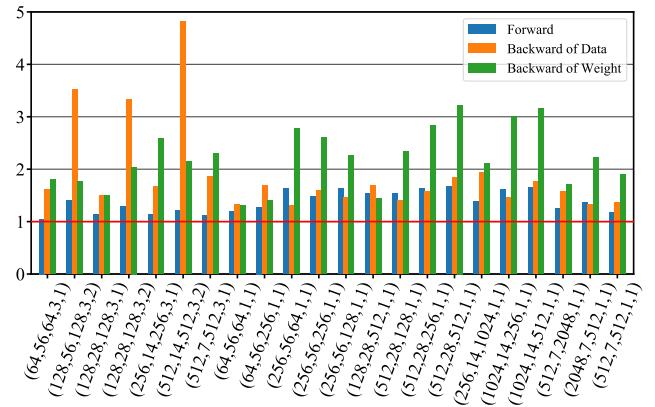


Figure 4: The speedup of INT8 over FP16 on convolution layer, where X-axis is (channel numbers, feature size, kernel number, kernel size, stride, batch size = 64), and Y-axis is speedup multiple.

Training Process	FP32	FP16	INT8
Forward (ms)	71.1	38.3	33.7
Backward (ms)	151.3	87.4	72.6
Iteration (ms)	237.8	136.4	115.6

Table 7: Running time of ResNet-50 on GeForce RTX 2080Ti (batch size = 64), including GQ and VQ overhead introduced in Algorithm 1.

parison. Our INT8 kernel is  $1.76\times$  faster than FP16 implementation with kernel size=1, and  $1.94\times$  faster with kernel size=3. Table 7 shows the average end-to-end time for one iteration of the training. We achieves 18% faster than FP16 implementation and more than 200% faster over FP32 version (UI8 only has 22% end-to-end speedup).

Considering FP16 quantization has been well studied, and supported by NVIDIA, it poses a great challenge to make INT8 quantization even faster in practice. To our best knowledge, we are the first to achieve actual speedup over FP16. We believe that if more engineering efforts are taken with INT8, it will bring more significant performance improvement.

## Conclusion

In this work, we well capture the gradient distribution as the key to achieve almost lossless quantization training for convolutional neural networks. In particular, we propose *Gradient Vectorized Quantization* to address multiple distributions along channel dimension, and *Magnitude-aware Clipping Strategy* to automatically update the optimal quantization parameters after introducing effect of gradient magnitude. Our approach has successfully trained on a lot of networks and deep learning tasks with negligible accuracy degradation, which sets a new state-of-the-art to the community. With the kernel level optimization, we further achieve more than 200% acceleration on modern Turing architecture.

## References

- Banner, R.; Hubara, I.; Hoffer, E.; and Soudry, D. 2018. Scalable methods for 8-bit training of neural networks. In *Advances in Neural Information Processing Systems (NIPS)*, 5145–5153.
- Cai, Z.; He, X.; Sun, J.; and Vasconcelos, N. 2017. Deep Learning with Low Precision by Half-Wave Gaussian Quantization. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 5406–5414.
- Chen, K.; Wang, J.; Pang, J.; Cao, Y.; Xiong, Y.; Li, X.; Sun, S.; Feng, W.; Liu, Z.; Xu, J.; Zhang, Z.; Cheng, D.; Zhu, C.; Cheng, T.; Zhao, Q.; Li, B.; Lu, X.; Zhu, R.; Wu, Y.; Dai, J.; Wang, J.; Shi, J.; Ouyang, W.; Loy, C. C.; and Lin, D. 2019. MMDetection: Open MMLab Detection Toolbox and Benchmark. *arXiv preprint arXiv:1906.07155*.
- Choi, J.; Venkataramani, S.; Srinivasan, V. V.; Gopalakrishnan, K.; Wang, Z.; and Chuang, P. 2019. Accurate and efficient 2-bit quantized neural networks. In *Machine Learning and Systems (MLSys)*.
- Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. Imagenet: A large-scale hierarchical image database. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 248–255.
- Everingham, M.; Van Gool, L.; Williams, C. K.; Winn, J.; and Zisserman, A. 2010. The pascal visual object classes (voc) challenge. *Int. J. Comput. Vis.* 88(2): 303–338.
- Gupta, S.; Agrawal, A.; Gopalakrishnan, K.; and Narayanan, P. 2015. Deep learning with limited numerical precision. In *International Conference on Machine Learning (ICML)*, 1737–1746.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 770–778.
- He, Z.; and Fan, D. 2019. Simultaneously optimizing weight and quantizer of ternary neural network using truncated gaussian approximation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 11438–11446.
- Hubara, I.; Courbariaux, M.; Soudry, D.; El-Yaniv, R.; and Bengio, Y. 2016. Binarized neural networks. In *Advances in Neural Information Processing Systems (NIPS)*, 4107–4115.
- 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 *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2704–2713.
- Jain, S. R.; Gural, A.; Wu, M.; and Dick, C. H. 2019. Trained quantization thresholds for accurate and efficient fixed-point inference of deep neural networks. In *Machine Learning and Systems (MLSys)*.
- Jung, S.; Son, C.; Lee, S.; Son, J.; Han, J.-J.; Kwak, Y.; Hwang, S. J.; and Choi, C. 2019. Learning to quantize deep networks by optimizing quantization intervals with task loss. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 4350–4359.
- Kay, W.; Carreira, J.; Simonyan, K.; Zhang, B.; Hillier, C.; Vijayanarasimhan, S.; Viola, F.; Green, T.; Back, T.; Natsev, P.; et al. 2017. The kinetics human action video dataset. *arXiv preprint arXiv:1705.06950*.
- Kolmogorov, A. 1933. Sulla determinazione empirica di una legge di distribuzione. *Inst. Ital. Attuari, Giorn.* 4: 83–91.
- Krizhevsky, A.; Sutskever, I.; and Hinton, G. E. 2012. Imagenet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems (NIPS)*, 1097–1105.
- Krizhevsky, A.; et al. 2009. Learning multiple layers of features from tiny images.
- Li, F.; Zhang, B.; and Liu, B. 2016. Ternary Weight Networks. In *NIPS Workshop on Efficient Methods for Deep Neural Networks*.
- Lin, T.-Y.; Goyal, P.; Girshick, R.; He, K.; and Dollár, P. 2017a. Focal loss for dense object detection. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2980–2988.
- Lin, T.-Y.; Maire, M.; Belongie, S.; Hays, J.; Perona, P.; Ramanan, D.; Dollár, P.; and Zitnick, C. L. 2014. Microsoft coco: Common objects in context. In *European Conference on Computer Vision (ECCV)*, 740–755.
- Lin, Y.; Han, S.; Mao, H.; Wang, Y.; and Dally, W. J. 2017b. Deep gradient compression: Reducing the communication bandwidth for distributed training. In *International Conference on Learning Representations (ICLR)*.
- Liu, Z.; Wu, B.; Luo, W.; Yang, X.; Liu, W.; and Cheng, K.-T. 2018. Bi-real net: Enhancing the performance of 1-bit cnns with improved representational capability and advanced training algorithm. In *European conference on computer vision (ECCV)*, 722–737.
- Loshchilov, I.; and Hutter, F. 2017. Sgdr: Stochastic gradient descent with warm restarts. In *International Conference on Learning Representations (ICLR)*.
- McCallum, W. G.; Hughes-Hallett, D.; Gleason, A. M.; Mumford, D.; Osgood, B. G.; Tecosky-Feldman, J.; Tucker, T. W.; Quinney, D.; and Lock, P. F. 2012. *Calculus: Multi-variable*. John Wiley.
- Nagel, M.; Baalen, M. v.; Blankevoort, T.; and Welling, M. 2019. Data-free quantization through weight equalization and bias correction. In *IEEE International Conference on Computer Vision (ICCV)*, 1325–1334.
- Patel, J. K.; and Read, C. B. 1996. *Handbook of the normal distribution*, volume 150. CRC Press.
- Ren, S.; He, K.; Girshick, R.; and Sun, J. 2015. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in Neural Information Processing Systems (NIPS)*, 91–99.
- Sandler, M.; Howard, A.; Zhu, M.; Zhmoginov, A.; and Chen, L.-C. 2018. Mobilenetv2: Inverted residuals and linear bottlenecks. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 4510–4520.



- Simonyan, K.; and Zisserman, A. 2015. Very deep convolutional networks for large-scale image recognition. In *International Conference on Learning Representations (ICLR)*.
- Song, L.; Pan, P.; Zhao, K.; Yang, H.; Chen, Y.; Zhang, Y.; Xu, Y.; and Jin, R. 2020. Large-Scale Training System for 100-Million Classification at Alibaba. In *ACM SIGKDD International Conference on Knowledge Discovery Data Mining (KDD)*, 2909–2930.
- Soomro, K.; Zamir, A. R.; and Shah, M. 2012. UCF101: A dataset of 101 human actions classes from videos in the wild. *arXiv preprint arXiv:1212.0402*.
- Strom, N. 2015. Scalable distributed DNN training using commodity GPU cloud computing. In *Conference of the International Speech Communication Association (INTER-SPEECH)*.
- Szegedy, C.; Vanhoucke, V.; Ioffe, S.; Shlens, J.; and Wojna, Z. 2016. Rethinking the inception architecture for computer vision. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2818–2826.
- Wang, N.; Choi, J.; Brand, D.; Chen, C.-Y.; and Gopalakrishnan, K. 2018a. Training deep neural networks with 8-bit floating point numbers. In *Advances in Neural Information Processing Systems (NIPS)*, 7675–7684.
- Wang, P.; Hu, Q.; Zhang, Y.; Zhang, C.; Liu, Y.; and Cheng, J. 2018b. Two-step quantization for low-bit neural networks. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 4376–4384.
- Wu, S.; Li, G.; Chen, F.; and Shi, L. 2018. Training and Inference with Integers in deep neural networks. In *International Conference on Learning Representations (ICLR)*.
- Yang, Y.; Deng, L.; Wu, S.; Yan, T.; Xie, Y.; and Li, G. 2020. Training high-performance and large-scale deep neural networks with full 8-bit integers. *Neural Networks* 125: 70–82.
- Zhang, D.; Yang, J.; Ye, D.; and Hua, G. 2018. Lq-nets: Learned quantization for highly accurate and compact deep neural networks. In *European Conference on Computer Vision (ECCV)*, 365–382.
- Zhang, J.; Pan, Y.; Yao, T.; Zhao, H.; and Mei, T. 2019. dabnn: A super fast inference framework for binary neural networks on arm devices. In *ACM International Conference on Multimedia*, 2272–2275.
- Zhang, X.; Liu, S.; Zhang, R.; Liu, C.; Huang, D.; Zhou, S.; Guo, J.; Guo, Q.; Du, Z.; Zhi, T.; et al. 2020. Fixed-point back-propagation training. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2330–2338.
- Zhou, S.; Wu, Y.; Ni, Z.; Zhou, X.; Wen, H.; and Zou, Y. 2016. Dorefa-net: Training low bitwidth convolutional neural networks with low bitwidth gradients. *arXiv preprint arXiv:1606.06160*.
- Zhu, F.; Gong, R.; Yu, F.; Liu, X.; Wang, Y.; Li, Z.; Yang, X.; and Yan, J. 2020. Towards unified int8 training for convolutional neural network. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 1969–1979.
- Zhuang, B.; Shen, C.; Tan, M.; Liu, L.; and Reid, I. 2018. Towards effective low-bitwidth convolutional neural networks. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 7920–7928.