

---

# Differentiable Dynamic Quantization with Mixed Precision and Adaptive Resolution

---

Zhang Zhaoyang<sup>1</sup> Shao Wenqi<sup>1</sup> Gu Jinwei<sup>2,3</sup> Wang Xiaogang<sup>1</sup> Luo Ping<sup>4</sup>

## Abstract

Model quantization is challenging due to many tedious hyper-parameters such as precision (bitwidth), dynamic range (minimum and maximum discrete values) and stepsize (interval between discrete values). Unlike prior arts that carefully tune these values, we present a fully differentiable approach to learn all of them, named Differentiable Dynamic Quantization (DDQ), which has several benefits. (1) DDQ is able to quantize challenging lightweight architectures like MobileNets, where different layers prefer different quantization parameters. (2) DDQ is hardware-friendly and can be easily implemented using low-precision matrix-vector multiplication, making it capable in many hardware such as ARM. (3) DDQ reduces training runtime by 25% compared to state-of-the-arts. Extensive experiments show that DDQ outperforms prior arts on many networks and benchmarks, especially when models are already efficient and compact. *e.g.* DDQ is the first approach that achieves lossless 4-bit quantization for MobileNetV2 on ImageNet.

## 1. Introduction

Deep Neural Networks (DNNs) have made significant progress in many applications. However, large memory and computations impede deployment of DNNs on portable devices. Model quantization (Courbariaux et al., 2015; 2016; Zhu et al., 2017) that discretizes the weights and activations of a DNN becomes an important topic, but it is challenging because of two aspects. Firstly, different network architectures allocate different memory and computational complexity in different layers, making quantization suboptimal when the quantization parameters such as bitwidth,

dynamic range and stepsize are frozen in each layer. Secondly, gradient-based training of low-bitwidth models is difficult (Bengio et al., 2013), because the gradient of previous quantization function may vanish, *i.e.* back-propagation through a quantized DNN may return zero gradients.

Previous approaches typically use the round operation with predefined quantization parameters, which can be summarized below. Let  $x$  and  $x_q$  be values before and after quantization, we have  $x_q = \text{sign}(x) \cdot d \cdot \mathcal{F}(\lfloor |x|/d + 0.5 \rfloor)$ , where  $|\cdot|$  denotes the absolute value,  $\text{sign}(\cdot)$  returns the sign of  $x$ ,  $d$  is the stepsize (*i.e.* the interval between two adjacent discrete values after quantization), and  $\lfloor \cdot \rfloor$  denotes the round function<sup>2</sup>. Moreover,  $\mathcal{F}(\cdot)$  is a function that maps a rounded value to a desired quantization level (*i.e.* a desired discrete value). For instance, the above equation would represent a uniform quantizer<sup>3</sup> when  $\mathcal{F}$  is an identity mapping, or it would represent a power-of-two (nonuniform) quantizer (Miyashita et al., 2016; Zhou et al., 2017; Liss et al., 2018; Zhang et al., 2018a; Li et al., 2020) when  $\mathcal{F}$  is a function of power of two.

Although using the round function is simple, model accuracy would drop when applying it on lightweight architectures such as MobileNets (Howard et al., 2019; Sandler et al., 2018) as observed by (Jain et al., 2019; Gong et al., 2019), because such models that are already efficient and compact do not have much room for quantization. As shown in Fig.1, small model is highly optimized for both efficiency and accuracy, making different layers prefer different quantization parameters. Applying predefined quantizers to them will result in large quantization errors (*i.e.*  $\|x - x_q\|_2$ ), which significantly decreases model accuracy compared to their full-precision counterparts even after finetuning the network weights.

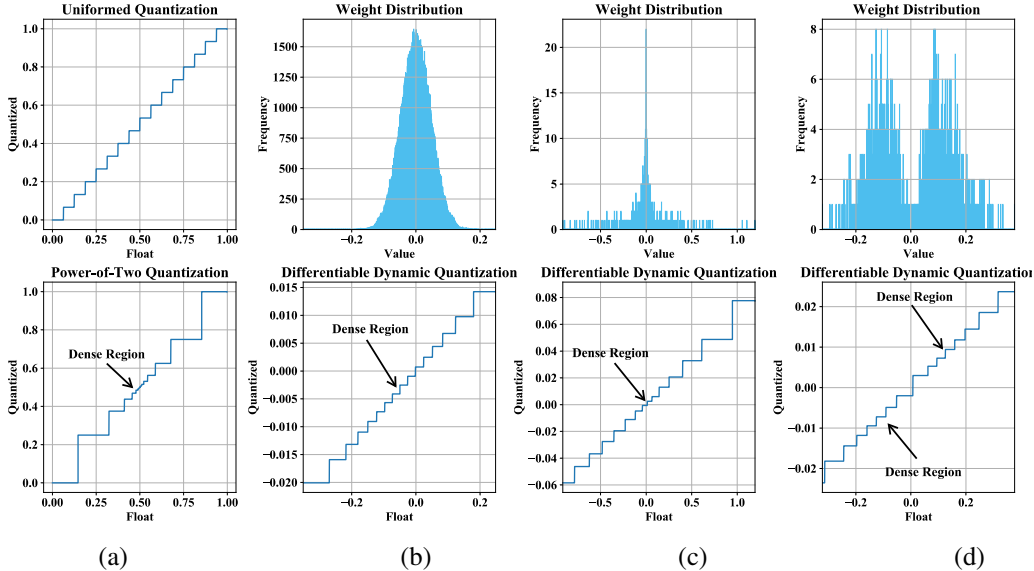
To improve model accuracy, recent quantizers reduce  $\|x - x_q\|_2$ . For example, TensorRT (Migacz, 2017), FAQ (McKinstry et al., 2018), PACT (Choi, 2018) and TQT (Jain et al., 2019) optimize an additional parameter that represents the

---

<sup>1</sup>The Chinese University of Hong Kong <sup>2</sup>SenseBrain, Ltd  
<sup>3</sup>Shanghai AI Lab <sup>4</sup>Hong Kong University. Correspondence to: Zhaoyang Zhang <zhaoyangzhang@link.cuhk.edu.hk>.

<sup>2</sup>The round function returns the closest discrete value given a continuous value.

<sup>3</sup>A uniform quantizer has uniform quantization levels, which means the stepsizes between any two adjacent discrete values are the same.



**Figure 1. Comparisons** between methods in 4-bit quantization. **(a)** compares quantization levels between uniform and nonuniform (power-of-two) (Miyashita et al., 2016; Zhou et al., 2017; Liss et al., 2018; Zhang et al., 2018a) quantizers, where  $x$ - and  $y$ -axis represent values before and after quantization respectively<sup>4</sup> (the float values are scaled between 0 and 1 for illustration). We highlight the dense region with higher “resolution” by arrows. We see that there is no dense region in uniform quantization (top), because the intervals between levels are the same, while a single dense region in power-of-two quantization (bottom). **(b)**, **(c)** and **(d)** show the distributions of network weights of different layers in a MobileNetv2 (Sandler et al., 2018) trained on ImageNet (Russakovsky et al., 2015), and the corresponding quantization levels learned by our DDQ. We see that different layers of MobileNetv2 prefer different parameters, *e.g.* weight distributions are Gaussian-like in (b), heavy-tailed in (c), and two-peak bell-shaped in (d). DDQ learns different quantization levels with different number of dense regions to capture these distributions.

dynamic range to calibrate the quantization levels, in order to better fit the distributions of the full-precision network. Besides, prior arts (Miyashita et al., 2016; Zhou et al., 2017; Liss et al., 2018; Zhang et al., 2018a; Li et al., 2020) also adopted non-uniform quantization levels with different step-sizes.

Despite the above works may reduce gap between values before and after quantization, they have different assumptions that may not be applied to efficient networks. For example, one key assumption is that the network weights would follow a bell-shaped distribution. However, this is not always plausible in many architectures such as MobileNets (Sandler et al., 2018; Howard et al., 2017; 2019), EfficientNets (Tan & Le, 2019), and ShuffleNets (Zhang et al., 2018b; Ma et al., 2018). For instance, Fig.1(b-d) plot different weight distributions of a MobileNetv2 (Sandler et al., 2018) trained on ImageNet, which is a representative efficient model on embedded devices. We find that they have irregular forms, especially when depthwise or group convolution (Xie et al., 2017; Huang et al., 2018; Zhang et al., 2017; Nagel et al., 2019; Park & Yoo, 2020) are used to improve computational efficiency. Although this problem has been identified in (Krishnamoorthi, 2018; Jain et al., 2019; Goncharenko et al., 2019), showing that per-channel scaling or bias correction

<sup>4</sup>For example, a continuous value is mapped to its nearest discrete value.

could compensate some of the problems, however, none of them could bridge the accuracy gap between the low-precision (*e.g.* 4-bit) and full-precision models. We provide full comparisons with previous works in Appendix B.

To address the above issues, we contribute by proposing Differentiable Dynamic Quantization (DDQ), to differentially learn quantization parameters for different network layers, including bitwidth, quantization level, and dynamic range. DDQ has appealing benefits that prior arts may not have. (1) DDQ is applicable in efficient and compact networks, where different layers are quantized by different parameters in order to achieve high accuracy. (2) DDQ is hardware-friendly and can be easily implemented using a low-precision matrix-vector multiplication (GEMM), making it capable in hardware such as ARM. Moreover, a matrix re-parameterization method is devised to reduce the matrix size from  $O(2^{2b})$  to  $O(\log 2^b)$ , where  $b$  is the number of bits. (3) Extensive experiments show that DDQ outperforms prior arts on many networks on ImageNet and CIFAR10, especially for the challenging small models *e.g.* MobileNets. To our knowledge, it is the first time to achieve lossless accuracy when quantizing MobileNet with 4 bits. For instance, MobileNetv2+DDQ achieves comparable top-1 accuracy on ImageNet, *i.e.* 71.9% *v.s.* 71.9% (full-precision), while ResNet18+DDQ improves the full-precision model, *i.e.* 71.3% *v.s.* 70.5%.

## 2. Our Approach

### 2.1. Preliminary and Notations

In network quantization, each continuous value  $x \in \mathbb{R}$  is discretized to  $x_q$ , which is an element from a set of discrete values. This set is denoted as a vector  $\mathbf{q} = [q_1, q_2, \dots, q_n]^T$  (termed quantization levels). We have  $x_q \in \mathbf{q}$  and  $n = 2^b$  where  $b$  is the bitwidth. Existing methods often represent  $\mathbf{q}$  using uniform or powers-of-two distribution introduced below.

**Uniform Quantization.** The quantization levels  $\mathbf{q}_u$  of a symmetric  $b$ -bit uniform quantizer (Uhlich et al., 2020) is

$$\mathbf{q}_u(\theta) = \left[ -1, \dots, \frac{-1}{2^{b-1}-1}, -0, +0, \frac{1}{2^{b-1}-1}, \dots, 1 \right]^T \times c + \bar{x}, \quad (1)$$

where  $\theta = \{b, c\}^T$  denotes a set of quantization parameters,  $b$  is the bitwidth,  $c$  is the clipping threshold, which represents a symmetric dynamic range<sup>6</sup>, and  $\bar{x}$  is a constant scalar (a bias) used to shift  $\mathbf{q}_u$ <sup>7</sup>. For example,  $\mathbf{q}_u(\theta)$  is shown in the upper plot of Fig.1(a) when  $c = 0.5$  and  $\bar{x} = 0.5$ . Although uniform quantization could be simple and effective, it assumes the weight distribution is uniform that is implausible in many recent DNNs.

**Powers-of-Two Quantization.** The quantization levels  $\mathbf{q}_p$  of a symmetric  $b$ -bit powers-of-two quantizer (Miyashita et al., 2016; Liss et al., 2018) is

$$\mathbf{q}_p(\theta) = \left[ -2^{-1}, \dots, -2^{-2^{b-1}+1}, -0, +0, 2^{-2^{b-1}+1}, \dots, 2^{-1} \right]^T \times c + \bar{x}. \quad (2)$$

As shown in the bottom plot of Fig.1(a) when  $c = 1$  and  $\bar{x} = 0.5$ ,  $\mathbf{q}_p(\theta)$  has a single dense region that may capture a single-peak bell-shaped weight distribution. Power-of-two quantization can also capture multiple-peak distribution by using the predefined additive scheme (Li et al., 2020).

In Eqn.(1-2), both uniform and power-of-two quantizers would fix  $\mathbf{q}$  and optimize  $\theta$ , which contains the clipping threshold  $c$  and the stepsize denoted as  $d = 1/(2^b - 1)$  (Uhlich et al., 2020). Although they learn the dynamic range and stepsize, they have an obvious drawback, that is, the predefined quantization levels cannot fit varied distributions of weights or activations for each layer during training.

### 2.2. Dynamic Differentiable Quantization (DDQ)

**Formulation for Arbitrary Quantizaion.** Instead of freezing the quantization levels, DDQ learns all quantization hyperparameters. Let  $Q(x; \theta)$  be a function with a set of parameters  $\theta$  and  $x_q = Q(x; \theta)$  turns a continuous value  $x$  into an element of  $\mathbf{q}$  denoted as  $x_q \in \mathbf{q}$ , where  $\mathbf{q}$  is initialized as a uniform quantizer and can be implemented in

<sup>6</sup>In Eqn.(1), the dynamic range is  $[-c, c]$ .

<sup>7</sup>Note that ‘0’ appears twice in order to assure that  $\mathbf{q}_u$  is of size  $2^b$ .

low-precision values according to hardware’s requirement. DDQ is formulated by low-precision matrix-vector product,

$$x_q = \mathbf{q}^T \frac{\mathbf{U}}{Z_U} \mathbf{x}_o, \text{ where } x_o^i = \begin{cases} 1 & \text{if } i = \operatorname{argmin}_j |\frac{1}{Z_U} (\mathbf{U}^T \mathbf{q})_j - x| \\ 0 & \text{otherwise} \end{cases}, \quad (3)$$

where  $x_o^i \in \mathbf{x}_o$ ,  $\mathbf{x}_o \in \{0, 1\}^{n \times 1}$  denotes a binary vector that has only one entry of ‘1’ while others are ‘0’, in order to select one quantization level from  $\mathbf{q}$  for the continuous value  $x$ . Note that we reparameterize  $\mathbf{q}$  by a trainable vector  $\tilde{\mathbf{q}}$ , such that

$$\mathbf{q} = R(\tilde{\mathbf{q}})(x_{max} - x_{min})/(2^b - 1) + x_{min}, \quad (4)$$

where  $R()$  denotes a uniform quantization function transforming  $\tilde{\mathbf{q}}$  to given  $b_q$  bits ( $b_q < b$ ). Eqn.(3) has parameters  $\theta = \{\tilde{\mathbf{q}}, \mathbf{U}\}$ , which are trainable by stochastic gradient descent (SGD), making DDQ automatically capture weight distributions of the full-precision models as shown in Fig.1(b-d). Here  $\mathbf{U} \in \{0, 1\}^{n \times n}$  is a binary block-diagonal matrix and  $Z_U$  is a constant normalizing factor used to average the discrete values in  $\mathbf{q}$  in order to learn bitwidth. Intuitively, different values of  $\mathbf{x}_o$ ,  $\mathbf{U}$  and  $\mathbf{q}$  make DDQ represent different quantization approaches as discussed below. To ease understanding, Fig.2(a) compares the computational graph of DDQ with the rounding-based methods. We see that DDQ learns the entire quantization levels instead of just the stepsize  $d$  as prior arts did.

**Discussions of Representation Capacity.** DDQ represents a wide range of quantization methods. For example, when  $\mathbf{q} = \mathbf{q}_u$  (Eqn.(1)),  $Z_U = 1$ , and  $\mathbf{U} = \mathbf{I}$  where  $\mathbf{I}$  is an identity matrix, Eqn.(3) represents an ordinary uniform quantizer. When  $\mathbf{q} = \mathbf{q}_p$  (Eqn.(2)),  $Z_U = 1$ , and  $\mathbf{U} = \mathbf{I}$ , Eqn.(3) becomes a power-of-two quantizer. When  $\mathbf{q}$  is learned, it represents arbitrary quantization levels with different dense regions.

Moreover, DDQ enables mixed precision training when  $\mathbf{U}$  is block-diagonal. For example, as shown in Fig.2(b), when  $\mathbf{q}$  has length of 8 entries (*i.e.* 3-bit),  $Z_U = \frac{1}{2}$ , and  $\mathbf{U} = \operatorname{Diag}(\mathbf{1}_{2 \times 2}, \dots, \mathbf{1}_{2 \times 2})$ , where  $\operatorname{Diag}(\cdot)$  returns a matrix with the desired diagonal blocks and its off-diagonal blocks are zeros and  $\mathbf{1}_{2 \times 2}$  denotes a 2-by-2 matrix of ones,  $\mathbf{U}$  enables Eqn.(3) to represent a 2-bit quantizer by averaging neighboring discrete values in  $\mathbf{q}$ . For another example, when  $\mathbf{U} = \operatorname{Diag}(\mathbf{1}_{4 \times 4}, \mathbf{1}_{4 \times 4})$  and  $Z_U = \frac{1}{4}$ , Eqn.(3) turns into a 1-bit quantizer. Besides, when  $\mathbf{x}_o$  is a soft one-hot vector with multiple non-zero entries, Eqn.(3) represents soft quantization that one continuous value can be mapped to multiple discrete values.

**Efficient Inference on Hardware.** DDQ is a unified quantizer that supports adaptive  $\mathbf{q}$  as well as predefined ones *e.g.* uniform and power-of-two. It is friendly to hardware with limited resources. As shown in Eqn.(3), DDQ reduces to a

uniform quantizer when  $q$  is uniform. In this case, DDQ can be efficiently computed by a rounding function as the step size is determined by  $U$  after training (*i.e.* don't have  $U$  and matrix-vector product when deploying in hardware like other uniform quantizers). In addition, DDQ with adaptive  $q$  can be implemented using low-precision general matrix multiply (GEMM). For example, let  $y$  be a neuron's activation,  $y = Q(w; \theta)x_q = \mathbf{q}^T \frac{U}{Z_U} \mathbf{w}_o x_q$ , where  $x_q$  is a discretized feature value,  $w$  is a continuous weight parameter to be quantized, and  $U$  and  $\mathbf{w}_o$  are binary matrix and one-hot vector of  $w$  respectively. To accelerate, we can calculate the major part  $\frac{U}{Z_U} \mathbf{w}_o x_q$  using low-precision GEMM first and then multiplying a short 1-d vector  $\mathbf{q}$ , which is shared for all convolutional weight parameters and can be float32, float16 or INT8 given specific hardware requirement. The latency in hardware is compared in later discussion.

### 2.3. Matrix Reparameterization of $U$

In Eqn.(3),  $U$  is a learnable matrix variable, which is challenging to optimize in two aspects. First, to make DDQ a valid quantizer,  $U$  should have binary block-diagonal structure, which is difficult to learn by using SGD. Second, the size of  $U$  (number of parameters) increases when the bitwidth increases *i.e.*  $2^{2b}$ . Therefore, rather than directly optimize  $U$  in the backward propagation using SGD, we explicitly construct  $U$  by composing a sequence of small matrices in the forward propagation (Luo et al., 2019).

**A Kronecker Composition for Quantization.** The matrix  $U$  can be reparameterized to reduce number of parameters from  $2^{2b}$  to  $\log 2^b$  to ease optimization. Let  $\{U_1, U_2, \dots, U_b\}$  denote a set of  $b$  small matrices of size 2-by-2,  $U$  can be constructed by  $U = U_1 \otimes U_2 \otimes \dots \otimes U_b$ , where  $\otimes$  denotes Kronecker product and each  $U_i$  ( $i = 1 \dots b$ ) is either a 2-by-2 identity matrix (denoted as  $I$ ) or an all-one matrix (denoted as  $\mathbf{1}_{2 \times 2}$ ), making  $U$  block-diagonal after composition. For instance, when  $b = 3$ ,  $U_1 = U_2 = I$  and  $U_3 = \mathbf{1}_{2 \times 2}$ , we have  $U = \text{Diag}(\mathbf{1}_{2 \times 2}, \dots, \mathbf{1}_{2 \times 2})$  and Eqn.(3) represents a 2-bit quantizer as shown in Fig.2(b).

To pursue a more parameter-saving composition, we further parameterize each  $U_i$  by using a single trainable variable. As shown in Fig.2(c), we have  $U_i = g_i I + (1 - g_i) \mathbf{1}_{2 \times 2}$ , where  $g_i = H(\hat{g}_i)$  and  $H(\cdot)$  is a Heaviside step function<sup>8</sup>, *i.e.*  $g_i = 1$  when  $\hat{g}_i \geq 0$ ; otherwise  $g_i = 0$ . Here  $\{g_i\}_{i=1}^b$  is a set of gating variables with binary values. Intuitively, each  $U_i$  switches between a 2-by-2 identity matrix and a 2-by-2 all-one matrix.

In other words,  $U$  can be constructed by applying a series of Kronecker products involving only  $\mathbf{1}_{2 \times 2}$  and  $I$ . Instead of updating the entire matrix  $U$ , it can be learned by only

<sup>8</sup>The Heaviside step function returns '0' for negative arguments and '1' for positive arguments.

a few variables  $\{\hat{g}_i\}_{i=1}^b$ , significantly reducing the number of parameters from  $2^b \times 2^b = 2^{2b}$  to  $b$ . In summary, the parameter size to learn  $U$  is merely the number of bits. With Kronecker composition, the quantization parameters of DDQ is  $\theta = \{g, \{\hat{g}_i\}_{i=1}^b\}$ , which could be different for different layers or kernels (*i.e.* layer-wise or channel-wise quantization) and the parameter size is negligible compared to the network weights, making different layers or kernels have different quantization levels and bitwidth.

**Discussions of Relationship between  $U$  and  $g$ .** Let  $g$  denote a vector of gates  $[g_1, \dots, g_b]^T$ . In general, different values of  $g$  represent different block-diagonal structures of  $U$  in two aspects. (1) **Permutation.** As shown in Fig.2(c),  $\{g_i\}_{i=1}^b$  should be permuted in an descending order by using a permutation matrix. Otherwise,  $U$  is not block-diagonal when  $g$  is not ordered, making DDQ an invalid quantizer. For example,  $g = [0, 1, 0]$  is not ordered compared to  $g = [1, 0, 0]$ . (2) **Sum of Gates.** Let  $s = \sum_{i=1}^b g_i$  be the sum of gates and  $0 \leq s \leq b$ . We see that  $U$  is a block-diagonal matrix with  $2^s$  diagonal blocks, implying that  $U^T \mathbf{q}$  has  $2^s$  different discrete values and represents a  $s$ -bit quantizer. For instance, as shown in Fig.2(b,c) when  $b = 3$ ,  $g = [1, 0, 0]^T$  and  $U = \text{Diag}(\mathbf{1}_{4 \times 4}, \mathbf{1}_{4 \times 4})$ , we have a  $s = 1 + 0 + 0 = 1$  bit quantizer. DDQ enables to regularize the value of  $s$  in each layer given memory constraint, such that optimal bitwidth can be assigned to different layers of a DNN.

## 3. Training DNN with DDQ

DDQ is used to train a DNN with mixed precision to satisfy memory constraints, which reduce the memory to store the network weights and activations, making a DNN appealing for deployment in embedded devices.

### 3.1. DNN with Memory Constraint

Consider a DNN with  $L$  layers trained using DDQ, the forward propagation of each layer can be written as

$$\mathbf{y}^l = F(Q(\mathbf{W}^l; \theta^l) * Q(\mathbf{y}^{l-1}) + Q(\mathbf{b}^l; \theta^l)), l = 1, 2, \dots, L \quad (5)$$

where  $*$  denotes convolution,  $\mathbf{y}^l$  and  $\mathbf{y}^{l-1}$  are the output and input of the  $l$ -th layer respectively,  $F$  is a non-linear activation function such as ReLU, and  $Q$  is the quantization function of DDQ. Let  $\mathbf{W}^l \in \mathbb{R}^{C_{out}^l \times C_{in}^l \times K^l \times K^l}$  and  $\mathbf{b}^l$  denote the convolutional kernel and bias vector (network weights), where  $C_{out}$ ,  $C_{in}$ , and  $K$  are the output and input channel size, and the kernel size respectively. Remind that in DDQ,  $\{g_i^l\}_{i=1}^b$  is a set of gates at the  $l$ -th layer and the bitwidth can be computed by  $s^l = \sum_{i=1}^b g_i^l$ . For example, the total memory footprint (denoted as  $\zeta$ ) can be computed



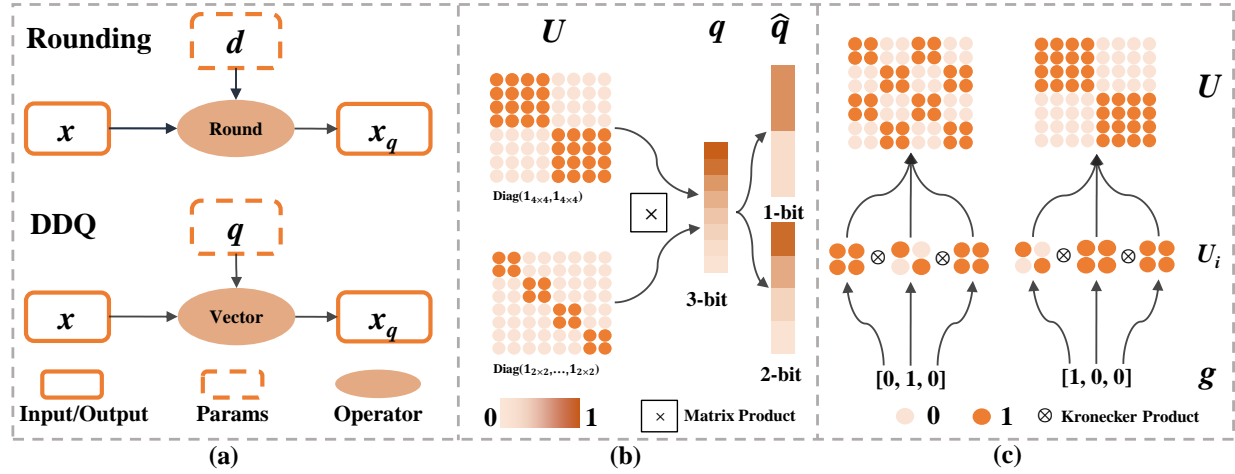


Figure 2. **Illustrations of DDQ.** (a) compares computations of DDQ with the round operator. Unlike rounding methods that only learn the stepsize  $d$ , DDQ treats  $q$  as trainable variable, learning arbitrary quantization levels. (b) illustrates that DDQ enables mixed-precision training by using different binary block-diagonal matrix  $U$ . The circles in light and dark indicate ‘0’ and ‘1’ respectively. For example, when  $q$  is of length 8 entries (*i.e.* 3-bit) and  $U = \text{Diag}(\mathbf{1}_{2 \times 2}, \dots, \mathbf{1}_{2 \times 2})$ , where  $\text{Diag}(\cdot)$  returns a matrix with the desired diagonal blocks while the off-diagonal blocks are zeros and  $\mathbf{1}_{2 \times 2}$  denotes a 2-by-2 matrix of ones, we have  $\hat{q} = U^T q$  that enables DDQ to represent a 2-bit quantizer by averaging neighboring discrete values in  $q$ . Another example is a 1-bit quantizer when  $U = \text{Diag}(\mathbf{1}_{4 \times 4}, \mathbf{1}_{4 \times 4})$ . (c) shows relationship between gating variables  $g = \{g_i\}_{i=1}^b$  and  $U$ . For example, when the entries of  $g = [1, 0, 0]$  are arranged in an descending order and let  $s = \sum_{i=1}^b g_i = 1 + 0 + 0 = 1$ ,  $U$  has  $2^s = 2^1 = 2$  number of all-one diagonal blocks. In such case, DDQ is a  $s = 1$  bit quantizer.

by

$$\zeta(s^1, \dots, s^L) = \sum_{l=1}^L C_{out}^l C_{in}^l (K^l)^2 2^{s^l}, \quad (6)$$

which represents the memory to store all network weights at the  $l$ -th layer when the bitwidth is  $s^l$ .

If the desired memory is  $\zeta(b^1, \dots, b^L)$ , we could use a weighted product to approximate the Pareto optimal solutions to train a  $b$ -bit DNN. The loss function is

$$\min_{\mathbf{w}^l, \theta^l} \mathcal{L}(\{\mathbf{W}^l\}_{l=1}^L, \{\theta^l\}_{l=1}^L) \cdot \left( \frac{\zeta(b^1, \dots, b^L)}{\zeta(s^1, \dots, s^L)} \right)^\alpha \quad (7)$$

s.t.  $\zeta(s^1, \dots, s^L) \leq \zeta(b^1, \dots, b^L)$ ,

where the loss  $\mathcal{L}(\cdot)$  is reweighted by the ratio between the desired memory and the practical memory similar to (Tan et al., 2019; Deb, 2014).  $\alpha$  is a hyper-parameter. We have  $\alpha = 0$  when the memory constraint is satisfied. Otherwise,  $\alpha < 0$  is used to penalize the memory consumption of the network.

### 3.2. Updating Quantization Parameters

All parameters of DDQ can be optimized by using SGD. This section derives their update rules. We omit the superscript ‘ $l$ ’ for simplicity.

**Gradients *w.r.t.*  $q$ .** To update  $q$ , we reparameterize  $q$  by a trainable vector  $\tilde{q}$ , such that  $q = R(\tilde{q})(x_{max} - x_{min}) / (2^{b-1} + x_{min})$ , in order to make each quantization level lies in  $[x_{min}, x_{max}]$ , where  $x_{max}$  and  $x_{min}$  are the maximum and minimum continuous values of a layer, and  $R(\cdot)$  denotes a uniform quantization function transforming  $\tilde{q}$  to given  $b_q$  bits ( $b_q < b$ ). Let  $x$  and  $x_q$  be two vectors stacking values before and after quantization respectively (*i.e.*  $x$  and  $x_q$ ), the gradient of the loss function with respect to each entry  $q_k$  of  $q$  is given by

$$\frac{\partial \mathcal{L}}{\partial q_k} = \sum_{i=1}^N \frac{\partial \mathcal{L}}{\partial x_q^i} \frac{\partial x_q^i}{\partial q_k} = \frac{1}{Z_U} \sum_{i \in \mathcal{S}_k} \frac{\partial \mathcal{L}}{\partial x_q^i}, \quad (8)$$

where  $x_q^i = Q(x^i; \theta)$  is the output of DDQ quantizer and  $\mathcal{S}_k$  represents a set of indexes of the values discretized to the corresponding quantization level  $q_k$ . In Eqn.(8), we see that the gradient with respect to the quantization level  $q_k$  is the summation of the gradients  $\frac{\partial \mathcal{L}}{\partial x_q^i}$ . In other words, the quantization level in denser regions would have larger gradients. The gradient *w.r.t.* gate variables  $\{g_i\}_{i=1}^b$  are discussed in Appendix A.

**Gradient Correction for  $x_q$ .** In order to reduce the quantization error  $\|x_q - x\|_2^2$ , a gradient correction term is proposed to regularize the gradient with respect to the quantized values,

$$\frac{\partial \mathcal{L}}{\partial x} \leftarrow \frac{\partial \mathcal{L}}{\partial x_q}, \quad \frac{\partial \mathcal{L}}{\partial x_q} \leftarrow \frac{\partial \mathcal{L}}{\partial x_q} + \lambda(x_q - x), \quad (9)$$




Figure 3. Learned quantization policy of each layer for ResNet18 and MobileNetV2 trained by DDQ on ImageNet. DDQ learns to allocate more bits to lower layers and depthwise layers of the networks.

where the first equation holds by applying STE. In Eqn.(9), we first assign the gradient of  $x_q$  to that of  $x$  and then add a correction term  $\lambda(x_q - x)$ . In this way, the corrected gradient can be back-propagated to the quantization parameters  $q$  and  $\{g_i\}_{i=1}^b$  in Eqn.(8), while not affecting the gradient of  $x$ . Intuitively, this gradient correction term is effective and can be deemed as the  $\ell_2$  penalty on  $\|x - x_q\|_2^2$ . Please note that this is not equivalent to simply impose a  $\ell_2$  regularization directly on the loss function, which would have no effect when STE is presented.

---

#### Algorithm 1 Training procedure of DDQ

---

**Input:** the full precision kernel  $W$  and bias kernel  $b$ , quantization parameter  $\theta = \{q, \{g_i\}_{i=1}^b\}$ , the target bitwidth of  $b_m$ , input activation  $y_{in}$ .

**Output:** the output activation  $y_{out}$

- 1: Apply DDQ to the kernel  $W$ , bias  $b$ , input activation  $y_{in}$  by Eqn.(3)
  - 2: Compute the output activation  $y_{out}$  by Eqn.(4)
  - 3: Compute the loss  $\mathcal{L}$  by Eqn.(6) and gradients  $\frac{\partial \mathcal{L}}{\partial y_{out}}$
  - 4: Compute the gradient of ordinary kernel weights and bias  $\frac{\partial \mathcal{L}}{\partial W}, \frac{\partial \mathcal{L}}{\partial m}$
  - 5: Applying gradient correction in Eqn.(9) to compute the gradient of parameters,  $\frac{\partial \mathcal{L}}{\partial q}, \frac{\partial \mathcal{L}}{\partial g_i}$  by Eqn.(7) and Eqn.(8)
  - 6: Update  $W, m, g_i$  and  $q$ .
- 

**Implementation Details.** Training DDQ can be simply implemented in existing platforms such as PyTorch and Tensorflow. The forward propagation only involves differentiable functions except the Heaviside step function. In practice, STE can be utilized to compute its gradients, *i.e.*  $\frac{\partial x_q}{\partial x} = 1_{\hat{q}_{min} \leq x \leq \hat{q}_{max}}$  and  $\frac{\partial g_k}{\partial \hat{g}_k} = 1_{|\hat{g}_k| \leq 1}$ , where  $\hat{q}_{min}$  and  $\hat{q}_{max}$  are minimum and maximum value in  $\hat{q} = U^T q$ . Algorithm. 1 provides detailed procedure. The codes will be released.

## 4. Experiments

We extensively compare DDQ with existing state-of-the-art methods and conduct multiple ablation studies on ImageNet (Russakovsky et al., 2015) and CIFAR dataset (Krizhevsky et al., 2009) (See Appendix D for eval-

uation on CIFAR). The reported validation accuracy is simulated with  $b_q = 8$ , if no other states.

### 4.1. Evaluation on ImageNet

**Comparisons with Existing Methods.** Table 1 compares DDQ with existing methods in terms of model size, bitwidth, and top1 accuracy on ImageNet using MobileNetV2 and ResNet18, which are two representative networks for portable devices. We see that DDQ outperforms recent state-of-the-art approaches by significant margins in different settings. For example, When trained for 30 epochs, MobileNetV2+DDQ yields 71.7% accuracy when quantizing weights and activations using 4 and 8 bit respectively, while achieving 71.5% when training with 4/4 bit. These results only drop 0.2% and 0.4% compared to the 32-bit full-precision model, outperforming all other quantizers, which may decrease performance a lot (2.4%~10.5%). For ResNet18, DDQ outperforms all methods even the full-precision model (*i.e.* 71.0% vs 70.5%). More importantly, DDQ is trained for 30 epochs, reducing the training time compared to most of the reported approaches that trained much longer (*i.e.* 90 or 120 epochs). Note that PROFIT (Park & Yoo, 2020) achieves 71.5% on MobileNet2 using a progressive training scheme (reducing bitwidth gradually from 8-bit to 5, 4-bit, 15 epochs each stage and 140 epochs in total.) This is quite similar to mixed-precision learning process of DDQ, in which bitwidth of each weight is initialized to maximum bits and learn to assign proper precision to each layer by decreasing the layerwise bitwidth. More details can see in Fig.4.

Fig.3 shows the converged bitwidth for each layer of MobileNet2 and ResNet18 trained with DDQ. We have two interesting findings. (1) Both networks tend to apply more bitwidth in lower layers, which have fewer parameters and thus being less regularized by the memory constraint. This allows to learn better feature representation, alleviating the performance drop. (2) As shown in the right hand side of Fig.3, we observe that depthwise convolution has larger bitwidth than the regular convolution. As found in (Jain et al., 2019), the depthwise convolution with irregular weight distributions is the main reason that makes quan-

|                                     | MobileNetV2     |            |                |                | ResNet18   |                |                |
|-------------------------------------|-----------------|------------|----------------|----------------|------------|----------------|----------------|
|                                     | Training Epochs | Model Size | Bitwidth (W/A) | Top-1 Accuracy | Model Size | Bitwidth (W/A) | Top-1 Accuracy |
| Full Precision                      | 120             | 13.2 MB    | 32             | 71.9           | 44.6 MB    | 32             | 70.5           |
| <b>DDQ (ours)</b>                   | 30              | 1.8 MB     | 4 / 8 (mixed)  | <b>71.7</b>    | 5.8 MB     | 4 / 8 (mixed)  | <b>71.0</b>    |
| <b>DDQ (ours)</b>                   | 90              | 1.8 MB     | 4 / 8 (mixed)  | <b>71.9</b>    | 5.8 MB     | 4 / 8 (mixed)  | <b>71.3</b>    |
| Deep Compression (Han et al., 2016) | -               | 1.8 MB     | 4 / 32         | 71.2           | -          | -              | -              |
| HMQ (Habi et al., 2020)             | 50              | 1.7 MB     | 4 / 32 (mixed) | 71.4           | -          | -              | -              |
| HAQ (Wang et al., 2019)             | 100             | 1.8 MB     | 4 / 32 (mixed) | 71.4           | -          | -              | -              |
| WPRN (Mishra et al., 2017)          | 100             | -          | -              | -              | 5.8 MB     | 4 / 8          | 66.4           |
| BCGD (Baskin et al., 2018)          | 80              | -          | -              | -              | 5.8 MB     | 4 / 8          | 68.9           |
| LQ-Net (Zhang et al., 2018a)        | 120             | -          | -              | -              | 5.8 MB     | 4 / 32         | 70.0           |
| <b>DDQ (ours)</b>                   | 90              | 1.8 MB     | 4 / 4 (mixed)  | <b>71.8</b>    | 5.8 MB     | 4 / 4 (mixed)  | <b>71.2</b>    |
| <b>DDQ (ours)</b>                   | 30              | 1.8 MB     | 4 / 4 (mixed)  | <b>71.5</b>    | 5.8 MB     | 4 / 4 (mixed)  | <b>71.0</b>    |
| <b>DDQ (ours)</b>                   | 90              | 1.8 MB     | 4 / 4 (fixed)  | <b>71.3</b>    | 5.8 MB     | 4 / 4 (fixed)  | <b>71.1</b>    |
| <b>DDQ (ours)</b>                   | 30              | 1.8 MB     | 4 / 4 (fixed)  | <b>70.7</b>    | 5.8 MB     | 4 / 4 (fixed)  | <b>70.7</b>    |
| PROFIT (Park & Yoo, 2020)           | 140             | 1.8 MB     | 4 / 4          | 71.5           | -          | -              | -              |
| SAT (Jin et al., 2019)              | 150             | 1.8 MB     | 4 / 4          | 71.1           | -          | -              | -              |
| HMQ (Habi et al., 2020)             | 50              | 1.7 MB     | 4 / 4 (mixed)  | 70.9           | -          | -              | -              |
| APOT (Li et al., 2020)              | 30              | 1.8 MB     | 4 / 4          | 69.7*          | -          | -              | -              |
| APOT (Li et al., 2020)              | 100             | 1.8 MB     | 4 / 4          | 71.0*          | 5.8 MB     | 4 / 4          | 70.7           |
| LSQ (Esser et al., 2020)            | 90              | 1.8 MB     | 4 / 4          | 70.6*          | 5.8 MB     | 4 / 4          | 71.1           |
| PACT (Choi, 2018)                   | 110             | 1.8 MB     | 4 / 4          | 61.4           | 5.6 MB     | 4 / 4          | 69.2           |
| DSQ (Gong et al., 2019)             | 90              | 1.8 MB     | 4 / 4          | 64.8           | 5.8 MB     | 4 / 4          | 69.6           |
| TQT (Jain et al., 2019)             | 50              | 1.8 MB     | 4 / 4          | 67.8           | 5.8 MB     | 4 / 4          | 69.5           |
| Uhlich et al. (Uhlich et al., 2020) | 50              | 1.6 MB     | 4 / 4 (mixed)  | 69.7           | 5.4 MB     | 4 / 4          | 70.1           |
| QIL (Jung et al., 2019)             | 90              | -          | -              | -              | 5.8 MB     | 4 / 4          | 70.1           |
| LQ-Net (Zhang et al., 2018a)        | 120             | -          | -              | -              | 5.8 MB     | 4 / 4          | 69.3           |
| NICE (Baskin et al., 2018)          | 120             | -          | -              | -              | 5.8 MB     | 4 / 4          | 69.8           |
| BCGD (Baskin et al., 2018)          | 80              | -          | -              | -              | 5.8 MB     | 4 / 4          | 67.4           |
| Dorefa-Net (Zhou et al., 2016)      | 110             | -          | -              | -              | 5.8 MB     | 4 / 4          | 68.1           |

Table 1. Comparisons between DDQ and state-of-the-art quantizers on ImageNet. “W/A” means bitwidth of weight and activation respectively. Mixed precision approaches are annotated as “mixed”. “-” denotes the absence of data in previous papers. We see that DDQ outperforms prior arts with much less training epochs. \* denotes our re-implemented results using public codes. Note that PROFIT (Park & Yoo, 2020) achieves 71.5% on MobileNetV2 using a progressive training scheme (reducing bitwidth gradually from 8-bit to 5, 4-bit, 15 epochs each stage and 140 epochs in total).




Figure 4. Evolution of bitwidth of top layers when training ResNet18. We can see that DDQ can learn to assign bitwidth to each layer under given memory footprint constraints. Full dynamics please see in Appendix. Quantizing MobileNet difficult. With mixed-precision training, DDQ allocates more bitwidth to depthwise convolution to alleviate this difficulty.

## 4.2. Ablation Study

**Ablation Study I: mixed versus fixed precision.** Mixed-precision quantization is new in the literature and proven

|                              | UQ   | PoT  | DDQ (fixed) | DDQ (mixed) |
|------------------------------|------|------|-------------|-------------|
| Maximum bitwidth             | 4    | 4    | 4           | 6 8         |
| Target bitwidth              | 4    | 4    | 4           | 4 4         |
| Weight memory footprint      | 1×   | 1×   | 1×          | 0.98× 1.03× |
| Top-1 Accuracy (MobileNetV2) | 65.2 | 68.6 | 70.7        | 71.2 71.5   |
| Weight memory footprint      | 1×   | 1×   | 1×          | 1.01× 0.96× |
| Top-1 Accuracy (ResNet18)    | 70.0 | 67.8 | 70.7        | 70.9 71.0   |

Table 2. Comparisons between PACT(Choi et al., 2018)+UQ, PACT(Choi et al., 2018)+PoT and DDQ on ImageNet. “DDQ (fixed)” and “DDQ (mixed)” indicate DDQ trained with fixed / mixed bitwidth. We see that DDQ+mixed surpasses all counterparts.

to be superior to their fixed bitwidth counterparts (Wang et al., 2019; Uhlich et al., 2020; Cai & Vasconcelos, 2020; Habi et al., 2020). DDQ is naturally used to perform mixed-precision training by a binary block-diagonal matrix  $U$ . In DDQ, each layer is quantized between 2-bit and a given maximum precision, which may be 4-bit / 6-bit / 8-bit. Items of gate  $\{\hat{g}_i\}_{i=1}^b$  are initialized all positively to  $1e - 8$ , which means  $U$  is identity matrix and precision of layers are initialized to their maximum values. We set target bitwidth as 4, constraining models to 4-bit memory footprint, and then jointly train  $\{\hat{g}_i\}_{i=1}^b$  with other parameters of corresponding model and quantizers. For memory constraints,  $\alpha$  is set to  $-0.02$  empirically. We use learning rates  $1e - 8$  towards  $\{\hat{g}_i\}_{i=1}^b$ , ensuring sufficient training when precision decreasing. Fig. 4 depicts the evolution of bitwidth for each layer when quantizing a 4-bit ResNet18 using DDQ with

|                              | Full Precision | UQ    | PoT   | DDQ(fix) | DDQ(fix) + GradCorrect |
|------------------------------|----------------|-------|-------|----------|------------------------|
| Bitwidth (W/A)               | 32 / 32        | 4 / 8 | 4 / 8 | 4 / 8    | 4 / 8                  |
| Top-1 Accuracy (MobileNetV2) | 71.9           | 67.1  | 69.2  | 71.2     | <b>71.6</b>            |
| Top-1 Accuracy (ResNet18)    | 70.5           | 70.6  | 70.8  | 70.8     | <b>70.9</b>            |
| Bitwidth (W/A)               | 32 / 32        | 4 / 4 | 4 / 4 | 4 / 4    | 4 / 4                  |
| Top-1 Accuracy (MobileNetV2) | 71.9           | 65.2  | 68.4  | 70.4     | <b>70.7</b>            |
| Top-1 Accuracy (ResNet18)    | 70.5           | 70.0  | 68.8  | 70.6     | <b>70.8</b>            |
| Bitwidth (W/A)               | 32 / 32        | 2 / 4 | 2 / 4 | 2 / 4    | 2 / 4                  |
| Top-1 Accuracy (MobileNetV2) | 71.9           | -     | -     | 60.1     | <b>63.7</b>            |
| Top-1 Accuracy (ResNet18)    | 70.5           | 66.5  | 63.8  | 67.4     | <b>68.5</b>            |
| Bitwidth (W/A)               | 32 / 32        | 2 / 2 | 2 / 2 | 2 / 2    | 2 / 2                  |
| Top-1 Accuracy (MobileNetV2) | 71.9           | -     | -     | 51.1     | <b>55.4</b>            |
| Top-1 Accuracy (ResNet18)    | 70.5           | 62.8  | 62.4  | 65.7     | <b>66.6</b>            |

Table 3. Ablation studies of adaptive resolution and gradient correction. “UQ” and “PoT” denote uniform and power-of-two quantization respectively. “DDQ(fix)+GradCorrect” refers to DDQ with gradient correction but fixed bitwidth. “-” denotes training diverged. “4/8” denotes training with 4-bit weights and 8-bit activations. Here we find that DDQ with gradient correction shows stable performances gains against DDQ (w/o gradient correction) and UQ/PoT baselines.


maximum bitwidth 8. As demonstrated, DDQ could learn to assign bitwidth to each layer, in a data-driven manner.

Table 2 compares DDQ trained using mixed precision to different fixed-precision quantization setups, including DDQ with fixed precision, uniform (UQ) and power-of-two (PoT) quantization by PACT (Choi et al., 2018) with gradient calibration (Jain et al., 2019; Esser et al., 2020; Jin et al., 2019; Bhalgat et al., 2020).


When the target bitwidth is 4, we see that DDQ trained with mixed precision significantly reduces accuracy drop of MobileNetV2 from 6.7% (e.g. PACT+UQ) to 0.4%. See Appendix C.2 for more details.

**Ablation Study II: adaptive resolution.** We evaluate the proposed adaptive resolution by training DDQ with homogeneous bitwidth (i.e. fixed  $U$ ) and only updating  $q$ . Table 3 shows performance of DNNs quantized with various quantization levels. We see that UQ and PoT incur a higher loss than DDQ, especially for MobileNetV2. We ascribe this drop to the irregular weight distribution as shown in Fig 1. Specially, when applying 2-bit quantization, DDQ still recovers certain accuracy compared to the full-precision model, while UQ and PoT may not converge. To our knowledge, DDQ is the first method to successfully quantize 2-bit MobileNet without using full precision in the activations.

**Ablation Study III: gradient correction.** We demonstrate how gradient correction improves DDQ. Fig 5(a) plots the training dynamics of layerwise quantization error i.e.  $\|W_q - W\|_2^2$ . We see that “DDQ+gradient correction” achieves low quantization error at each layer (average error is 0.35), indicating that the quantized values well approximate their continuous values. Fig 5(b) visualizes the trained quantization levels. DDQ trained with gradient correction would capture the distribution to the original weights, thus reducing the quantization error. As shown in Table.3, training DDQ with gradient correction could apparently reduce accuracy drop from quantization.



(a) Training dynamics of quantization error.



(b) Learned quantization level for channels.

Figure 5. Training dynamics of quantization error in MobileNetV2. (a) compares the quantization errors of PACT+UQ, PACT+PoT, and DDQ with/without gradient correction. DDQ with gradient correction shows stable convergence and lower quantization errors than counterparts. (b) compares the converged quantization levels of DDQ for each channel with/without gradient correction, and dense regions are marked by arrows. Here we can also see that quantization levels learned by DDQ with gradient correction could fit original data distribution better.

### 4.3. Evaluation on Mobile Devices

In Table 4. we further evaluate DDQ on mobile platform to trade off of accuracy and latency. With implementation stated in Appendix C.1, DDQ achieves over 3% accuracy gains compared to UQ baseline. Moreover, in contrast to FP model, DDQ runs with about 40% less latency with just a small accuracy drop (<0.3%). Note that all tests here are



| Methods                     | bitwidth(w/a) | Mixed-precision | Latency(ms) | Top-1(%) |
|-----------------------------|---------------|-----------------|-------------|----------|
| FP                          | 32/32         |                 | 7.8         | 71.9     |
| UQ                          | 4/8           |                 | 3.9         | 67.1     |
| DDQ (fixed) <sup>1</sup>    | 4/8           |                 | 4.5         | 71.7     |
| DDQ (mixed UQ) <sup>2</sup> | 4/8           | ✓               | 4.1         | 70.8     |
| DDQ <sup>3</sup>            | 4/8           | ✓               | 5.1         | 71.9     |

Table 4. Comparison of Quantized MobileNetv2 runing on mobile DSPs. <sup>1</sup> Fixed precision DDQ. <sup>2</sup> Mixed precision DDQ with uniform quantizer constraints. <sup>3</sup> Original DDQ. “w/a” means the bitwidth for network weights and activations respectively.

running under INT8 simulation due to the support of the platform. We believe the acceleration ratio can be larger in the future when deploying DDQ on more compatible hardwares.

## 5. Conclusion

This paper introduced a differentiable dynamic quantization (DDQ), a versatile and powerful algorithm for training low-bit neural network, by automatically learning arbitrary quantization policies such as quantization levels and bitwidth. DDQ represents a wide range of quantizers. DDQ did well in the challenging MobileNet by significantly reducing quantization errors compared to prior work. We also show DDQ can learn to assign bitwidth for each layer of a DNN under desired memory constraints. Unlike recent methods that may use reinforcement learning(Wang et al., 2019; Yazdanbakhsh et al., 2018), DDQ doesn’t require multiple epochs of retraining, but still yield better performance compared to existing approaches.

## 6. Acknowledgement

This work is supported in part by Centre for Perceptual and Interactive Intelligence Limited, in part by the General Research Fund through the Research Grants Council of Hong Kong under Grants (Nos. 14202217, 14203118, 14208619), in part by Research Impact Fund Grant No. R5001-18. This work is also partially supported by the General Research Fund of HK No.27208720.

## References

- Baskin, C., Liss, N., Chai, Y., Zheltonozhskii, E., Schwartz, E., Giryas, R., Mendelson, A., and Bronstein, A. M. Nice: Noise injection and clamping estimation for neural network quantization. *arXiv preprint arXiv:1810.00162*, 2018.
- Bengio, Y., Léonard, N., and Courville, A. Estimating or propagating gradients through stochastic neurons for conditional computation. *arXiv preprint arXiv:1308.3432*, 2013.
- Bhalgat, Y., Lee, J., Nagel, M., Blankevoort, T., and Kwak, N. Lsq+: Improving low-bit quantization through learnable offsets and better initialization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pp. 696–697, 2020.
- Cai, Z. and Vasconcelos, N. Rethinking differentiable search for mixed-precision neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 2349–2358, 2020.
- Choi, J. Pact: Parameterized clipping activation for quantized neural networks. *arXiv: Computer Vision and Pattern Recognition*, 2018.
- Choi, J., Chuang, P. I.-J., Wang, Z., Venkataramani, S., Srinivasan, V., and Gopalakrishnan, K. Bridging the accuracy gap for 2-bit quantized neural networks (qnn). *arXiv preprint arXiv:1807.06964*, 2018.
- Courbariaux, M., Bengio, Y., and David, J.-P. Binaryconnect: Training deep neural networks with binary weights during propagations. In *Advances in neural information processing systems*, pp. 3123–3131, 2015.
- Courbariaux, M., Hubara, I., Soudry, D., El-Yaniv, R., and Bengio, Y. Binarized neural networks: Training deep neural networks with weights and activations constrained to+ 1 or-1. *arXiv preprint arXiv:1602.02830*, 2016.
- Deb, K. Multi-objective optimization. In *Search methodologies*, pp. 403–449. Springer, 2014.
- Esser, S. K., McKinstry, J. L., Bablani, D., Appuswamy, R., and Modha, D. S. Learned step size quantization. 2020.
- Goncharenko, A., Denisov, A., Alyamkin, S., and Terentev, E. Fast adjustable threshold for uniform neural network quantization. *International Journal of Computer and Information Engineering*, 13(9):499–503, 2019.
- Gong, R., Liu, X., Jiang, S., Li, T., Hu, P., Lin, J., Yu, F., and Yan, J. Differentiable soft quantization: Bridging full-precision and low-bit neural networks. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 4852–4861, 2019.

- Habi, H. V., Jennings, R. H., and Netzer, A. Hmq: Hardware friendly mixed precision quantization block for cnns. *arXiv preprint arXiv:2007.09952*, 2020.
- Han, S., Mao, H., and Dally, W. J. Deep compression: Compressing deep neural networks with pruning, trained quantization and huffman coding. *International Conference on Learning Representations*, 2016.
- Howard, A., Sandler, M., Chu, G., Chen, L.-C., Chen, B., Tan, M., Wang, W., Zhu, Y., Pang, R., Vasudevan, V., et al. Searching for mobilenetv3. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 1314–1324, 2019.
- Howard, A. G., Zhu, M., Chen, B., Kalenichenko, D., Wang, W., Weyand, T., Andreetto, M., and Adam, H. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861*, 2017.
- Huang, G., Liu, S., Van der Maaten, L., and Weinberger, K. Q. Condensenet: An efficient densenet using learned group convolutions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 2752–2761, 2018.
- Jain, S. R., Gural, A., Wu, M., and Dick, C. H. Trained quantization thresholds for accurate and efficient fixed-point inference of deep neural networks. *arXiv: Computer Vision and Pattern Recognition*, 2019.
- Jin, Q., Yang, L., and Liao, Z. Towards efficient training for neural network quantization. *arXiv preprint arXiv:1912.10207*, 2019.
- Jung, S., Son, C., Lee, S., Son, J., Han, J.-J., Kwak, Y., Hwang, S. J., and Choi, C. Learning to quantize deep networks by optimizing quantization intervals with task loss. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 4350–4359, 2019.
- Krishnamoorthi, R. Quantizing deep convolutional networks for efficient inference: A whitepaper. *arXiv preprint arXiv:1806.08342*, 2018.
- Krizhevsky, A., Hinton, G., et al. Learning multiple layers of features from tiny images. 2009.
- Li, Y., Dong, X., and Wang, W. Additive powers-of-two quantization: A non-uniform discretization for neural networks. 2020.
- Liss, N., Baskin, C., Mendelson, A., Bronstein, A. M., and Giryes, R. Efficient non-uniform quantizer for quantized neural network targeting reconfigurable hardware. *arXiv preprint arXiv:1811.10869*, 2018.
- Luo, P., Zhanglin, P., Wenqi, S., Ruimao, Z., Jiamin, R., and Lingyun, W. Differentiable dynamic normalization for learning deep representation. In *International Conference on Machine Learning*, pp. 4203–4211, 2019.
- Ma, N., Zhang, X., Zheng, H.-T., and Sun, J. Shufflenet v2: Practical guidelines for efficient cnn architecture design. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 116–131, 2018.
- McKinstry, J. L., Esser, S. K., Appuswamy, R., Bablani, D., Arthur, J. V., Yildiz, I. B., and Modha, D. S. Discovering low-precision networks close to full-precision networks for efficient embedded inference. *arXiv preprint arXiv:1809.04191*, 2018.
- Migacz, S. 8-bit inference with tensorsrt. In *GPU technology conference*, volume 2, pp. 5, 2017.
- Mishra, A., Nurvitadhi, E., Cook, J. J., and Marr, D. Wrpn: wide reduced-precision networks. *arXiv preprint arXiv:1709.01134*, 2017.
- Miyashita, D., Lee, E. H., and Murmann, B. Convolutional neural networks using logarithmic data representation. *arXiv preprint arXiv:1603.01025*, 2016.
- Nagel, M., Baalen, M. v., Blankevoort, T., and Welling, M. Data-free quantization through weight equalization and bias correction. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 1325–1334, 2019.
- Park, E. and Yoo, S. Profit: A novel training method for sub-4-bit mobilenet models. *arXiv preprint arXiv:2008.04693*, 2020.
- Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., et al. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115(3): 211–252, 2015.
- Sandler, M., Howard, A., Zhu, M., Zhmoginov, A., and Chen, L.-C. Mobilenetv2: Inverted residuals and linear bottlenecks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 4510–4520, 2018.
- Tan, M. and Le, Q. Efficientnet: Rethinking model scaling for convolutional neural networks. In *International Conference on Machine Learning*, pp. 6105–6114. PMLR, 2019.
- Tan, M., Chen, B., Pang, R., Vasudevan, V., Sandler, M., Howard, A., and Le, Q. V. Mnasnet: Platform-aware neural architecture search for mobile. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 2820–2828, 2019.

- Uhlich, S., Mauch, L., Cardinaux, F., Yoshiyama, K., Garcia, J. A., Tiedemann, S., Kemp, T., and Nakamura, A. Mixed precision dnns: All you need is a good parametrization. In *International Conference on Learning Representations*, 2020.
- Wang, K., Liu, Z., Lin, Y., Lin, J., and Han, S. Haq: Hardware-aware automated quantization with mixed precision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 8612–8620, 2019.
- Xie, S., Girshick, R., Dollár, P., Tu, Z., and He, K. Aggregated residual transformations for deep neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 1492–1500, 2017.
- Yazdanbakhsh, A., Elthakeb, A. T., Pilligundla, P., Mireshghallah, F., and Esmailzadeh, H. Releq: An automatic reinforcement learning approach for deep quantization of neural networks. 2018.
- Zhang, D., Yang, J., Ye, D., and Hua, G. Lq-nets: Learned quantization for highly accurate and compact deep neural networks. In *Proceedings of the European conference on computer vision (ECCV)*, pp. 365–382, 2018a.
- Zhang, T., Qi, G.-J., Xiao, B., and Wang, J. Interleaved group convolutions. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 4373–4382, 2017.
- Zhang, X., Zhou, X., Lin, M., and Sun, J. Shufflenet: An extremely efficient convolutional neural network for mobile devices. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 6848–6856, 2018b.
- Zhou, A., Yao, A., Guo, Y., Xu, L., and Chen, Y. Incremental network quantization: Towards lossless cnns with low-precision weights. *International Conference on Learning Representations*, 2017.
- Zhou, S., Wu, Y., Ni, Z., Zhou, X., Wen, H., and Zou, Y. Dorefa-net: Training low bitwidth convolutional neural networks with low bitwidth gradients. *arXiv preprint arXiv:1606.06160*, 2016.
- Zhu, C., Han, S., Mao, H., and Dally, W. J. Trained ternary quantization. *International Conference on Learning Representations*, 2017.