

# Any-Precision Deep Neural Networks\*

Haichao Yu<sup>1</sup>, Haoxiang Li<sup>2</sup>, Humphrey Shi<sup>1,3</sup>, Thomas S. Huang<sup>1</sup>, Gang Hua<sup>2</sup>

<sup>1</sup>UIUC, <sup>2</sup>Wormpex AI Research, <sup>3</sup>University of Oregon

{haichao3, hshi10, t-huang1}@illinois.edu, haoxiang.li@bianlifeng.com, ganghua@gmail.com

## Abstract

We present any-precision deep neural networks (DNNs), which are trained with a new method that allows the learned DNNs to be flexible in numerical precision during inference. The same model in runtime can be flexibly and directly set to different bit-widths, by truncating the least significant bits, to support dynamic speed and accuracy trade-off. When all layers are set to low-bits, we show that the model achieved accuracy comparable to dedicated models trained at the same precision. This nice property facilitates flexible deployment of deep learning models in real-world applications, where in practice trade-offs between model accuracy and runtime efficiency are often sought. Previous literature presents solutions to train models at each individual fixed efficiency/accuracy trade-off point. But how to produce a model flexible in runtime precision is largely unexplored. When the demand of efficiency/accuracy trade-off varies from time to time or even dynamically changes in runtime, it is infeasible to re-train models accordingly, and the storage budget may forbid keeping multiple models. Our proposed framework achieves this flexibility without performance degradation. More importantly, we demonstrate that this achievement is agnostic to model architectures and applicable to multiple vision tasks. Our code is released at <https://github.com/SHI-Labs/Any-Precision-DNNs>.

## Introduction

While state-of-the-art deep learning models can achieve very high accuracy on various benchmarks, runtime cost is another crucial factor to consider in practice. In general, the capacity of a deep learning model is positively correlated with its complexity. As a result, accurate models mostly run slower, consume more power, and have larger memory footprint as well as model size. In practice, it is inevitable to balance efficiency and accuracy to get a good trade-off when deploying any deep learning models.

To alleviate this issue, a number of approaches have been proposed to address it from different perspectives. We observe active researches (Liu et al. 2018a; Cai, Zhu, and Han 2018; Chen et al. 2019) in looking for more efficient

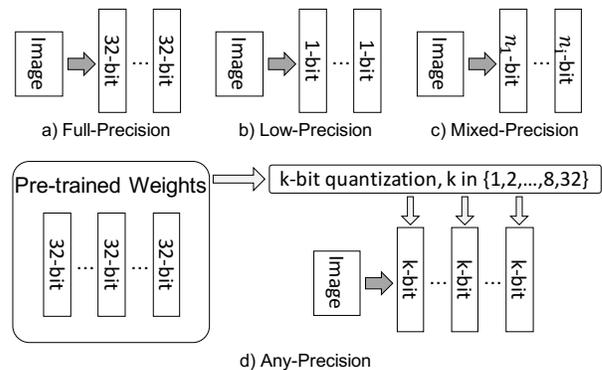


Figure 1: Illustrations of deep neural networks (DNN) in different numerical precisions: a) weights and activations of typical full-precision DNN are in 32-bit floating values; b) binary DNN with 1-bit weights and activations as an example low-precision DNN; c) different layers in mixed-precision DNN can be in arbitrary bit-width and  $n_i$  is fixed after training; d) the proposed any-precision DNN can have pre-trained weights in full-precision while in runtime the weights and activations can be quantized into arbitrary bit-width  $k$ .

deep neural network architectures to support practical usage (Howard et al. 2019; Yu and Huang 2019; Tan and Le 2019). People also consider to adaptively modify general deep learning model inference to dynamically determine the execution during the feed-forward pass to save some computation at the cost of potential accuracy drop (Figurnov et al. 2017; Teerapittayanon, McDanel, and Kung 2016; Wu et al. 2018; Veit and Belongie 2018).

Besides these explorations, another important line of research proposes a low-level solution to use less bits to represent deep learning model and its runtime data to achieve largely reduced runtime cost. It has been shown in various literatures that full-precision is over-abundant in many applications that we can use 8-bit or even 4-bit models without obvious performance degradation.

Some previous works went further in this direction. For example, BNN, XNOR-Net, and others (Courbariaux et al. 2016; Rastegari et al. 2016; Zhou et al. 2016) are proposed to use as low as 1-bit for both the weights and acti-

\*Correspondence to H. Shi and G. Hua. This work was done during Yu’s internship at Wormpex AI Research. Copyright © 2021, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

vations of the deep neural networks to reduce power-usage, memory-footprint, running time, and model size. However, ultra low-precision models always observe obvious accuracy drop (Courbariaux et al. 2016). While many methods have been proposed to improve accuracy of the low-precision models, so far we see no silver bullet. Stepping back from uniformly ultra-low precision models, mixed-precision models have been proposed to serve as a better trade-off (Wang et al. 2019; Dong et al. 2019). Effective ways have been found to train accurate models with some layers processing in ultra-low precision and some layers in high precision. We illustrate these different paradigms in Figure 1.a-c.

When we look at this spectrum of deep learning models in terms of its numerical precision, from full-precision at one end to low-precision at the other, and mixed-precision in between, we have to admit that efficiency/accuracy trade-off always exists in reality and to deploy a model in a specific application scenario we have to find the right trade-off point. Previous methods can provide a specific operating point but what if we demand flexibility as well? It would be a highly favorable property if we can dynamically change the efficiency/accuracy trade-off point given a single model. Preferably, we want to be able to adjust the model, without the need of re-training or re-calibration, to run in high accuracy mode when resources are sufficient and switch to low accuracy mode when resources are limited.

In this paper, we propose a method to train deep learning models to be flexible in numerical precision, namely any-precision deep neural networks. After training, we can freely quantize the model layers into various precision levels, without fine-tuning or calibration and without any data. We illustrate this in Figure 1.d. When running in low-precision, full-precision or other precision levels in between, it achieves comparable accuracy to models specifically trained under the matched settings. Furthermore, given fixed computational budget, it can potentially find better operating point than one trained rigorously.

To summarize, our contributions are:

- We introduce the concept of any-precision DNN. In runtime we can quantize its layers into different bit-widths. Its accuracy changes smoothly with respect to its precision level without drastic performance degradation;
- We propose a novel model-agnostic method to train any-precision DNN and validate its effectiveness over multiple vision tasks, with multiple widely used benchmarks, and with multiple neural network architectures;
- We demonstrate the proposed training framework trains better low-bit models with knowledge distillation.

## Related Work

**Low-Precision Deep Neural Networks.** Recent progresses in deep learning inference hardware motivate the research of using low-bit integer instead of float-point values to represent network weights and activations. Binarized Neural Networks (Courbariaux et al. 2016) and XNOR-Net (Rastegari et al. 2016) are early works in this direction to use only 1-bit to represent the weights and activations

in DNNs. When training these 1-bit networks, a float-point value copy of the parameters are maintained under the hood to calculate approximated gradients. Usually a *sign* function is used to quantize the float-point value copy to binary value in the feed-forward pass. Using only 1-bit numerical precision leads to obvious drop in accuracy in most scenarios, Zhou et al. (Zhou et al. 2016) proposed DoReFa-Net to specifically train arbitrary bitwidth in weights, activations, and gradients. Since gradients are also in low-bits, proper implementation could accelerate both the forward and backward passes.

One of the essential problem in learning low-precision DNNs is the quantization operator. Quantization of the real-value parameters in the feed-forward pass and approximation of the gradients through the quantization operator in the backward pass heavily influence the final model accuracy. For example, the *sign* function adopted in Binarized NN (Courbariaux et al. 2016) discards the value distribution variations across layers and hurt the performance. In XNOR-Net (Rastegari et al. 2016), a scaling factor is added to each layer to minimize the information loss. Choi et al. (Choi et al. 2018) proposed a parameterized clipping activation for quantization to support arbitrary bits quantization of activation. Zhang et al. (Zhang et al. 2018) and Jung et al. (Jung et al. 2019) pointed out that having a uniform quantization pattern across layers is suboptimal and propose a learnable quantizer for each layer to improve the model accuracy.

In the backward pass, most prior works use the Straight-Through Estimator (STE) (Bengio, Léonard, and Courville 2013) to approximate the gradients over the quantizers. Cai et al. (Cai et al. 2017) proposed to use a half-wave gaussian quantization operator to replace the *sign* function for better learning efficiency and a piece-wise continuous function in the backpropagation step to alleviate the gradient mismatch issue in the prior design. Liu et al. (Liu et al. 2018b) also attacked the gradient mismatch problem by introducing a piecewise polynomial function to approximate the *sign* function. Another interesting recent work from Ding et al. (Ding et al. 2019) addressed this problem by introducing a new loss function over the value distribution of layer activations.

Besides the performance gap to full-precision model, training binary networks have been reportedly to be unstable. Tang et al. (Tang, Hua, and Wang 2017) carefully analyzed the training process and concluded that using PReLU (He et al. 2015) activation function, a low learning rate, and the bipolar regularization on weights could lead to a more stable training process with better optimum. Zhuang et al. (Zhuang et al. 2018) looked at the overall training strategy and propose a progressive training process. They suggested to first train the net with quantized weights and then quantized activations, first train with high-precision and then low-precision, and jointly train the low-bit model with the full-precision one.

A similar joint training strategy has been observed to be effective in this work as well. Since our work is along an orthogonal direction of low-precision DNNs training and design, our method can be complementary to train better and flexible DNNs.

**Post-training Quantization.** Quantization of a pre-trained model with fine-tuning or calibration on a dataset is another related research topic in the area. Although methods in this area are working on a different problem from ours, we partially share the motivation to have the flexibility of quantization control in the runtime. Without special treatment, many models collapse even in 8-bit precision in post-training quantization. One recent work from Nagel *et al.* (Nagel et al. 2019) identified two issues leading to the large accuracy drop, the large variation in the weight ranges across channels and biased output errors due to quantization errors affecting following layers. With their method, they are able to alleviate the bias and equalize the weight ranges by rescaling and reparameterization. In this paper, the model we produced can be readily quantized into lower precision without further process.

In the research area of deep neural networks architecture search, the slimmable neural networks by Yu *et al.* (Yu et al. 2018) is related to ours in terms of methodology. They presented method to train a single neural network with adjustable number of channels in each layer at runtime. Their exploration is limited to the search space of network architecture instead of weights.

## Any-Precision Deep Neural Networks

### Overview

Neural networks are generally constructed layer by layer. We denote input to the  $i$ -th layer in a neural network model as  $\mathbf{x}_i$ , the weights of the layer as  $\mathbf{w}_i$  and the biases as  $\mathbf{b}_i$ . The output from this layer can be calculated as

$$\mathbf{y}_i = \mathcal{F}(\mathbf{x}_i | \mathbf{w}_i, \mathbf{b}_i). \quad (1)$$

Without loss of generality, we take one channel in a fully-connected layer as a concrete example in the following description and drop the subscript  $i$  for simplicity, i.e.,

$$\mathbf{y} = \mathbf{w} \cdot \mathbf{x} + b, \quad (2)$$

where  $\mathbf{y}, \mathbf{w} \in \mathbb{R}^D$  and  $b$  is a scalar.

For better computation efficiency, we would like to avoid the float-value dot product of  $D$ -dimensional vectors. Instead we use  $N$ -bit fixed-point integers to represent the weights as  $\mathbf{w}_Q$  and input activations as  $\mathbf{x}_Q$ . Hereafter, we assume  $\mathbf{w}_Q$  and  $\mathbf{x}_Q$  are stored as signed integers in its bit-wise format. Note that in some related works (Courbariaux et al. 2016), elements of  $\mathbf{w}_Q$  and  $\mathbf{x}_Q$  could be represented as vectors of  $\{-1, 1\}$  and the conversion between these two formats are trivial. With  $N$ -bit integers weights and activations, as discussed in prior arts (Zhou et al. 2016; Rastegari et al. 2016), the computation can be accelerated by leveraging bit-wise operations (*and*, *xnor*, *bit-count*), or even dedicated DNN hardware.

Early works (Tang, Hua, and Wang 2017) show that by adding a layer-wise real-value scaling factor  $s$  could largely help reduce the output range variation and hence achieve better model accuracy. Since the scaling factor is shared across channels within the same layer, the computational cost is fractional. Following this setting, with the quantized weights and inputs, we have

$$\mathbf{y}' = s * (\mathbf{w}_Q \cdot \mathbf{x}_Q) + b. \quad (3)$$

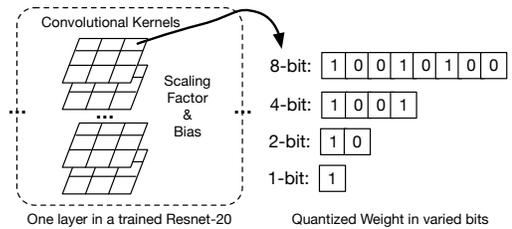


Figure 2: Quantization of a kernel weight in the trained model into different precision levels: since we follow an uniform quantization pattern, when representing weight values in signed integers, the quantization can be implemented as simple bit-shift.

The activations  $\mathbf{y}'$  are then quantized into  $N$ -bit fixed-point integers as the input to the next layer.

### Inference

We will discuss our quantization functions in details in the next section. Here we describe the runtime of a trained any-precision DNN.

Once training is finished, we can keep the weights at a higher precision level for storage, for example, at 8-bit. As shown in Figure 2, we can simply quantize the weights into lower bit-width by bit-shifting. We experimentally observe that with the proposed training framework, the model accuracy changes smoothly and consistently on-par or even out-perform dedicated models trained at the same bit-width.

### Training

A number of quantization functions have been proposed in the literature for weights and activations respectively. Given a pre-trained DNN model, one can quantize its weights into low-bit and apply certain quantization function to activations accordingly. However, when the number of bits gets smaller, the accuracy quickly drops due to the rough approximation in weights and large variations in activations. The most widely adopted framework to obtain low-bit model is quantization-aware training. The proposed method follows the quantization-aware training framework.

We take the same fully-connected layer as an example. In training, we maintain the float-point value weights  $\mathbf{w}$  for the actual layer weights  $\mathbf{w}_Q$ . In the feed-forward pass, given input  $\mathbf{x}_Q$ , we follow Equation 3 to compute the raw output  $\mathbf{y}'$ . Prior arts show the importance of the batch normalization (BN) (Ioffe and Szegedy 2015) layer in low-precision DNN training and we follow accordingly.  $\mathbf{y}'$  is then passed into a BN layer and then quantized into  $\mathbf{y}^Q$  as the input to the next layer.

**Weights.** We use a uniform quantization strategy similar to Zhou *et al.* (Zhou et al. 2016) with a scaling factor to approximate the weights. Given the floating point weight  $\mathbf{w}$ , we first apply the *tanh* function to normalize it into  $[-1, 1]$  and then transform it into  $\mathbf{w}' \in [0, 1]$ , i.e.,

$$\mathbf{w}' = \frac{\tanh(\mathbf{w})}{2\max(|\tanh(\mathbf{w})|)} + 0.5. \quad (4)$$

We then quantize normalized value into  $N$ -bit integers  $\mathbf{w}_Q'$  and scaling factor  $s$ , where

$$\mathbf{w}_Q' = \text{INT}(\text{round}(\mathbf{w}' * \text{MAX}_N)) \text{ and } s' = 1 / \text{MAX}_N. \quad (5)$$

Hereafter  $\text{MAX}_N$  denotes the upper-bound of  $N$ -bit integer and  $\text{INT}(\cdot)$  converts a floating point value into an integer.

Finally the values are re-mapped back to approximate the range of floating point values to obtain

$$\mathbf{w}_Q = 2 * \mathbf{w}_Q' - 1 \text{ and } s = \mathbb{E}(|\mathbf{w}|) / \text{MAX}_N, \quad (6)$$

where  $\mathbb{E}$  is the mean of absolute value of all floating-valued weights in the same layer. Eventually, we approximate  $\mathbf{w}$  with  $s * \mathbf{w}_Q$  and execute the feed-forward pass with the quantized weights as shown in Equation 3, the scaling factor can be applied after the dot-product of all integers vectors.

In the backward pass, gradients are computed with respect to the underlying float-value variable  $\mathbf{w}$  and updates are applied to  $\mathbf{w}$  as well. In this way, the relatively unreliable and nuance signals would be accumulated gradually and hence this will stabilize the overall training process. Since not all operations involved are nice smooth functions to support back-propagation, we use the straight through estimator (STE) (Bengio, Léonard, and Courville 2013) to approximate the gradients. For example, the *round* operation in Equation 5 has zero derivative almost everywhere. With STE, we assign  $\partial \text{round}(x) / \partial x := 1$ .

**Activations.** For activation quantization in the feed-forward pass, we obtain the  $N$ -bit fixed-point representation by first clipping the value to be within  $[0, 1]$  and then

$$\begin{aligned} \mathbf{y}_c' &= \text{clip}(\mathbf{y}', 0, 1), \\ \mathbf{y}_Q &= \text{INT}(\text{round}(\mathbf{y}_c' * \text{MAX}_N)) * \frac{1}{\text{MAX}_N}, \end{aligned} \quad (7)$$

In practice, we only calculate the integer part as  $\mathbf{y}_Q$  and absorb the constant scaling factor into the persistent network parameters in the next layer.

Let  $L$  denote the final loss function and the gradient with respect to the activation  $\mathbf{y}_Q$  is then approximated to be

$$\frac{\partial L}{\partial \mathbf{y}_Q} \approx \frac{\partial L}{\partial \mathbf{y}_c'}, \quad (8)$$

where

$$\frac{\partial L}{\partial \mathbf{y}_c'} = \begin{cases} \frac{\partial L}{\partial \mathbf{y}'}, & \text{if } 0 \leq \mathbf{y}' \leq 1, \\ 0, & \text{otherwise.} \end{cases} \quad (9)$$

The gradient of the *round* function is approximated with STE to be 1.

**Dynamic Model-wise Quantization.** In prior low-precision models, the bit-width  $N$  is fixed during the training process. In runtime, if we alter  $N$  the model accuracy drops drastically. To encourage flexibility in the produced model, here we propose to dynamically change  $N$  within the training stage to align the training and inference process. However, the distribution of activations varies under different bit-width  $N$ , especially when  $N$  is small

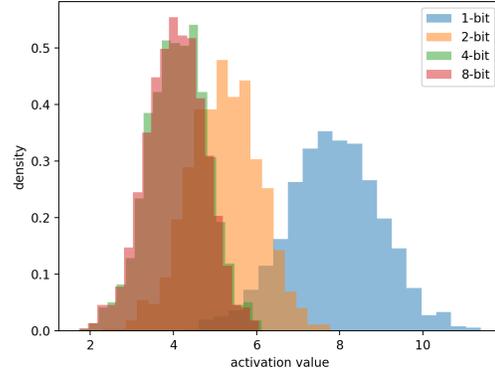


Figure 3: Activation distributions under different bit-widths for weights and inputs: we randomly generate a single-channel fully-connected layer and 1000 16-dimensional inputs; we then quantize the weights and inputs into 1, 2, 4, 8 bits respectively and summarize the distributions of activations under different bit-widths; as observed in the figure, the 1-bit quantization leads to significant distribution shift compared to 8-bit model and the discrepancy under 2-bit is also obvious.

(e.g., 1-bit), as shown in Figure 3. As a result, without special treatment, the dynamically changed  $N$  creates conflicts in learning the model that it fails to converge in our experiments.

One of the widely adopted technique to adjust internal feature/activation distribution is Batch Normalization (BatchNorm) (Ioffe and Szegedy 2015). It works by normalizing layer output across batch dimension as following

$$\hat{x}_i = \gamma \frac{x_i - \mu}{\sqrt{\sigma^2 + \epsilon}} + \beta, \quad i = 1..B, \quad (10)$$

where  $B$  is the batch size,  $i$  denotes the index within current batch,  $\epsilon$  is a small value added to avoid numerical issue.  $\mu$  and  $\sigma^2$  are mean and variance respectively defined as

$$\mu = \sum_{i=1}^B x_i \text{ and } \sigma^2 = \frac{1}{B} \sum_{i=1}^B (x_i - \mu)^2. \quad (11)$$

During training, BatchNorm layer keeps calculating running averages for  $\mu$  and  $\sigma^2$ , i.e.,

$$\mu = \lambda \mu + (1 - \lambda) \mu_t \text{ and } \sigma^2 = \lambda \sigma^2 + (1 - \lambda) \sigma_t^2, \quad (12)$$

where  $\mu_t$  and  $\sigma_t^2$  are the values before the current update, the decay rate  $\lambda$  is a hyper-parameter set a-prior. But even with the BatchNorm layer, dynamically changed  $N$  will lead to failure of convergence in training due to the value distribution variations shown in the toy example in Figure 3.

In our proposed framework, we adopted **dynamically changed BatchNorm layer** to work with different  $N$  in training. More specifically, assume we have a list of bit-width candidates  $\{n_k\}_{k=1}^K$ , we keep  $|K|$  copies of BatchNorm layer parameters and internal states  $\Phi_{k=1}^K$ . When the current training iteration works with  $N = n_k$ , we reset the BatchNorm layers with data from  $\Phi_k$  to use and update the corresponded copy.

---

**Algorithm 1** Training of the proposed any-precision DNN

---

**Require:** Given candidate bit-widths  $P \leftarrow \{n_k\}_{k=1}^K$

- 1: Initialize the model  $\mathcal{M}$  with floating-value parameters
- 2: Initialize  $K$  BatchNorm layers:  $\Phi_{k=1}^K$
- 3: **for**  $t = 1, \dots, T_{iters}$  **do**
- 4:   Sample data batch  $(x, y)$  from train set  $D_{train}$
- 5:   **for**  $n_p$  in  $P$  **do**
- 6:     Set quantization bit-width  $N \leftarrow n_p$
- 7:     Feed-forward pass:  $y_{n_p} \leftarrow \mathcal{M}(x)$
- 8:     Set BatchNorm layers:  $\mathcal{M}.\text{replace}(\Phi_p)$
- 9:      $L \leftarrow L + \text{loss}(y_{n_p}, y)$
- 10:   **end for**
- 11:   Back-propagate to update network parameters
- 12: **end for**

---

Similar technique has been adopted by Yu *et al.* (Yu and Huang 2019) when dealing with varied network architectures. Parameters of all BatchNorm layers are kept after training and used in inference. Note that compared with the total number of network parameters, the additional amount from BatchNorm layers is negligible. We summarize the proposed method in Algorithm 1. With the proposed algorithm, we can train DNN being flexible for runtime bit-width adjustment.

Another optional component in our method is **adding knowledge distillation (Hinton, Vinyals, and Dean 2015) in training**. Knowledge distillation works by matching the outputs of two networks. In training a network, we can use a more complicated model or an ensemble of models to produce soft targets by adjusting the temperature of the final softmax layer and then use the soft targets to guide the network learning.

In our framework, we apply this idea by generating soft targets from a high-precision model. More specifically, in each training iteration, we first set the quantization bit-width to the highest candidate  $n_K$  and run feed-forward pass to obtain soft targets  $y_{soft}$ . Later, instead of accumulating cross-entropy loss for each precision candidate, we use KL divergence of the model prediction and  $y_{soft}$  as the loss. In our experiments, we observe that in general knowledge distillation leads to better performance at low-bit precision levels.

## Experiments

We first validate our method with several network architectures and datasets on image classification task. These networks include a 8-layer CNN (named Model C in (Zhou *et al.* 2016)), AlexNet (Krizhevsky, Sutskever, and Hinton 2012), MobileNetV2 (Sandler *et al.* 2018), Resnet-8, Resnet-18, Resnet-20 and Resnet-50 (He *et al.* 2016). The datasets include Cifar-10 (Krizhevsky, Hinton *et al.* 2009), Street View House Numbers (SVHN) (Netzer *et al.* 2011), and ImageNet (Deng *et al.* 2009). We also evaluate our method on the image segmentation task to demonstrate its generalization.

## Implementation Details

We implement the whole framework in PyTorch (Paszke *et al.* 2017). On Cifar-10, we train AlexNet, MobileNetV2 and Resnet-20 models for 400 epochs with initial learning rate 0.001 and decayed by 0.1 at epochs {150, 250, 350}. On SVHN, the 8-layer CNN named CNN-8 and Resnet-8 models are trained for 100 epochs with initial learning rate 0.001 and decayed by 0.1 at epochs {50, 75, 90}. We combine the training and extra training data on SVHN as our training dataset. All models on Cifar-10 and SVHN are optimized with the Adam optimizer (Kingma and Ba 2014) without weight decay. On ImageNet, we train Resnet-18 and Resnet-50 dedicated models for 120 epochs with initial learning rate 0.1 decayed by 0.1 at epochs {30, 60, 85, 95, 105} with SGD optimizer. For any-precision model, we trained 80 epochs with initial learning 0.3 decayed by 0.1 at epochs {45, 60, 70}.

For all models, following Zhou *et al.* (Zhou *et al.* 2016) we keep first and last layer real-valued. In training, we train the networks with bit-width candidates {1, 2, 4, 8, 32}. Note that when the bit-width is set to 32, it is a full-precision model that we use floating-valued weights and activations. In testing, we evaluate the model at each bit-width in the list respectively. By default, we recursively add knowledge distillation (KD) in training. Concretely, we use the full-precision model to get soft targets as supervision for the 8-bit model, the ones from the 8-bit model for the 4-bit model, so on and so forth.

## Comparison to Dedicated Models

We compare our method to very competitive baseline models at each precision level. For each bit-width we tested, we dedicatedly train a low-precision model following the same training pipeline with fixed bit-width for weights and activations. We compare the accuracy we obtained from our dedicated low-bit models to other recent works in this field to make sure the baseline models are strong. For example, on Cifar-10, our 1-bit baseline achieved an accuracy of 92.07% while the recent work from Ding *et al.* (Ding *et al.* 2019) reported 89.90%.

Our results are summarized in Table 1. We observed that binarizing depth-wise convolution layers in MobileNetV2 may lead to training divergence, which is verified by Hai *et al.* (Phan *et al.* 2020). Therefore, we replace depth-wise convolutions with group convolutions in MobileNetV2 models. As shown in the table, on all three datasets, the proposed any-precision DNN generally achieved comparable performance to the competitive dedicated models. At the same time, our model is more compact. As shown in Table 2, compared with five individual models, our any-precision model (104MB) saved more than 50% parameters (220MB).

## Post-Training Quantization Methods

We compare our method with three alternative post-training quantization methods. We experiment with Resnet-50 on ImageNet.

The first naive baseline directly quantizes dedicated models with bit-shifting. In other words, to obtain an  $(n - 1)$ -bit

Datasets	Models	1 bit		2 bit		4 bit		8 bit		FP32	
		Dedi.	Ours								
Cifar-10	Resnet-20	92.07	92.15	93.55	93.97	93.71	93.95	93.66	93.80	94.08	93.98
	AlexNet	92.56	93.00	94.06	94.08	94.02	94.26	93.82	94.24	93.74	94.22
	MobileNetV2	80.17	80.06	89.69	90.74	92.01	92.27	92.44	92.32	94.01	93.84
SVHN	Resnet-8	92.94	91.65	95.91	94.78	95.15	95.46	94.64	95.39	94.60	95.36
	CNN-8	90.94	88.21	96.45	94.94	97.04	96.19	97.04	96.22	97.10	96.29
ImageNet	Resnet-18	55.06	54.62	63.65	64.19	68.15	67.96	68.48	68.04	69.27	68.16
	Resnet-50	61.08	63.18	69.44	73.24	71.24	74.75	74.71	74.91	75.95	74.96

Table 1: Comparison of the proposed any-precision DNN to dedicated models: the proposed method achieved the strong baseline accuracy in most cases, even occasionally outperforms the baselines in low-precision. We hypothesize that the gain is mainly from the knowledge distillation from high-precision models in training. Dedi.: Dedicated models.

	Shared	1	2	4	8	32	Total
Ded	0	22.3	24.9	30.1	40.4	102.4	220.1
Ours	102.0	0.4	0.4	0.4	0.4	0.4	104.0

Table 2: Resnet-50 model size comparison (in MegaByte): the proposed model only brings little overhead upon size of a FP32 model but achieves the flexibility what could take 5 independent models instead. Ded: Dedicated models.

Models \ Runtime bit	1	2	4	8	FP32
BS	0.100	0.136	42.70	74.68	75.95
BS+BC	0.106	0.352	57.12	73.73	75.95
ACIQ	0.116	0.324	71.64 <sup>1</sup>	75.83	75.95
Ours	63.18	73.24	74.75	74.91	74.96

Table 3: Comparison to other post-training quantization methods: All models are Resnet-50 trained on ImageNet. When bit-width drops from their original training setting, our method consistently outperform them. BS: Bit-Shifting from FP32 model. BC: BatchNorm Calibration from FP32 model.

model from a trained  $n$ -bit model, as what is done with any-precision DNN shown in Figure 2, we simply drop the least-significant bit of all weights. With no surprise, this strategy fails dramatically on challenging large-scale benchmark as shown in Table 3.

The second strategy follows the same bit-shifting to drop bit with an added BatchNorm calibration process. In the calibration process, BatchNorm statistics will be re-calculated by feed-forwarding a number of training samples. As shown in Table 3, the BatchNorm calibration helped in 4-bit but still failed in 1,2-bit settings.

With the proposed method, we can leverage this post-training calibration technique to fill-in the gaps of training candidate bit-width list, i.e., after training for 1,2,4,8,32-bits precision levels, we can further calibrate the model under the remaining 3,5,6,7-bit settings to get the missed copies of BatchNorm layer parameters. So that, in runtime, we can freely choose any precision level from 1 to 8 bits.

In addition, we compare with a recently proposed method

<sup>1</sup>We used the paper’s official code <https://github.com/submission2019/cnn-quantization>, but we cannot reproduce the claimed accuracy 73.8% on 4-bit model.

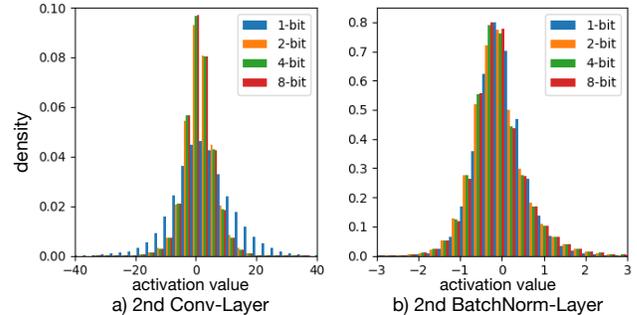


Figure 4: Activation value distributions of several layers in an any-precision AlexNet: low-bit quantization leads to value distribution different from others after convolutional layers but accordingly changed BatchNorm layer could rectify the mis-match.

Train \ Test	1	2	3	4	5	6	7	8
1,2,4,8	91.80	93.48	93.22	93.70	93.58	93.52	93.53	93.61
1,8	91.95	89.57	92.97	93.23	93.35	93.34	93.36	93.39
2,8	10.00	93.58	93.50	93.83	93.91	93.90	94.00	93.93
4,8	10.00	72.78	93.22	93.80	93.75	93.79	93.81	93.79

Table 4: Classification accuracy of Resnet-20 with different bit-width combinations in training on Cifar-10.

ACIQ (Banner, Nahshan, and Soudry 2019), which introduced analytical weight clipping, adaptive bit allocation and bias correction for post-training quantization. As shown in Table 3, our method achieved much better accuracy under low bit-width settings.

### Dynamically Changed BatchNorm Layers

To understand how the dynamically changed BatchNorm layers help in our framework, we visualize the activation value distributions of several layers of an any-precision AlexNet. More specifically, we look at how activation value distribution changes from the 2nd convolutional layers and the BatchNorm layers after them when the runtime precision level is set to 1,2,4,8-bit respectively. As shown in Figure 4, when running at 1-bit precision, the activation distribution is obviously off from others after the convolutional layers; the followed BatchNorm layer rectifies the distributions; then the next convolutional layer would create this dis-

Models \ Bits		1	2	4	8	FP32
Resnet-20 on Cifar-10	w/o KD	91.67	93.72	93.83	93.94	93.82
	KD	92.15	93.97	93.95	93.80	93.98
Resnet-50 on ImageNet	w/o KD	62.42	72.88	74.43	74.68	74.98
	KD	63.18	73.24	74.75	74.91	74.96

Table 5: Impact of knowledge distillation on classification accuracy.

Models \ Bits		1	2	4	8	FP32
Dedicated	mIoU	59.5	71.5	74.1	75.6	76.2
	Acc.	89.7	93.5	94.2	94.5	94.6
Ours	mIoU	61.3	72.1	74.3	75.4	75.9
	Acc.	90.6	93.7	94.3	94.4	94.6

Table 6: Segmentation performance comparison on Pascal VOC 2012 segmentation dataset under mIoU and top-1 accuracy.

tribution variation again. It is very clear that by keeping multiple copies of the BatchNorm layer parameters for different bit-widths, we can minimize input variations to the convolutional layers and hence have the same set of convolutional layer parameters to support any-precision in runtime.

## Ablation Studies

**Candidate bit-width List.** We study how the candidate bit-width list used during training the any-precision DNN influences the testing performance on other bit-widths.

Table 4 shows testing accuracy of models trained under different bit-width combinations. We observe that training with more candidate bit-width generally leads to better generalization to the others and the candidate bit-width list is better to cover the extreme cases in runtime. For example, the 1,8-bits combination performs more stable across different runtime bit-width compared with 2,8-bits and 4,8-bits combination. Since better coverage in training takes longer for the model to converge, this observation can guide the bit-width selection under limited training resources.

**Knowledge Distillation.** We study the influence of knowledge distillation during any-precision training. The first case is *w/o KD*, *i.e.*, no KD is used as shown in Algorithm 1). The second case is *KD*, *i.e.*, the highest bit-width outputs are supervised by groundtruth, and then every other bit-width outputs are supervised by the output logits from the nearest superior bit-width.

In Table 5, we observed on both Resnet-20 and Resnet-50, KD outperforms the dedicated models in low-bit settings. The hypothesis is that for lower-bit models, soft logits from higher-bit models instead of groundtruth labels regularize the training better.

## Application to Semantic Segmentation

To demonstrate generalization of our method to other tasks, we apply the any-precision scheme to semantic segmentation. We train Deeplab V3 (Chen et al. 2017) with Resnet-50 on Pascal VOC 2012 segmentation dataset (Everingham et al. 2015) with SBD dataset (Hariharan et al. 2011) as

groundtruth augmentation. We use the publicly available PyTorch training codebase<sup>2</sup>. All models are pretrained from the corresponding classification models on ImageNet. We do not use any post-processing for segmentation.

We summarize our results in Table 6 under two evaluation metrics mIoU and top-1 accuracy. Our low-bit models perform better than the dedicated counterparts. In 8,32-bit settings, our models achieve comparable results. Similar to the results on the classification tasks, we attribute this improvement to the joint optimization and knowledge distillation.

## Discussion

We study how the multiple bit-width joint training influences parameter learning. With STE, the parameter being updated in training is essentially the FP32 weight  $w$ . The gradient on  $w$  comes from losses from different bit-widths. The necessary condition for the joint training to work is that the gradients from different bit-widths are consistent with each other.

Inspired by (Zhao et al. 2018), we analyze gradient consistency between bit-widths by computing Update Compliance Average (UCA), which is defined as the average of cosine similarity of gradients from two bit-widths over multiple training steps. During joint training, we observed that different bit-widths share very large UCA ( $[0.9, 1]$ ), indicating consistent gradient directions and thus training convergence. We also found that neighboring bit-widths share larger UCA than the others, *e.g.*, UCA between 1 and 2 bits (0.929) are larger than 1 and the other bits (0.909 between 1 and 32 bits). This motivated us to employ recursive knowledge distillation in joint training.

As indicated by UCA, we also observed a small gradient direction gap. We think this makes our model naturally more robust to adversarial attacks than a single dedicated model as a by-product from joint training. By switching to different bit-widths, any-precision model under one bit-width can show robustness to adversarial attacks targeting at another bit-width. We experimented with FGSM attacking method (Goodfellow, Shlens, and Szegedy 2015). Any-precision models show improved defensive performance compared with dedicated ones, *e.g.*, classification accuracy of FP32 model is improved by more than 10%. We believe our any-precision scheme can serve as an add-on to other defensive methods against adversarial attacks.

## Conclusion

In this paper, we introduce any-precision DNN to address the practical efficiency/accuracy trade-off dilemma from a new perspective. Instead of seeking for a better operating point, we enable runtime adjustment of model precision-level to support flexible efficiency/accuracy trade-off without additional storage or computation cost. The model can be stored at 8-bit or higher and run in lower bit-width specified at runtime. Our flexible model achieves comparable accuracy to dedicatedly trained low-precision models and surpasses other post-training quantization methods.

<sup>2</sup><https://github.com/kazuto1011/deeplab-pytorch>

## References

- Banner, R.; Nahshan, Y.; and Soudry, D. 2019. Post training 4-bit quantization of convolutional networks for rapid-deployment. In *Advances in Neural Information Processing Systems*, 7948–7956.
- Bengio, Y.; Léonard, N.; and Courville, A. 2013. Estimating or propagating gradients through stochastic neurons for conditional computation. *arXiv preprint arXiv:1308.3432*.
- Cai, H.; Zhu, L.; and Han, S. 2018. Proxylessnas: Direct neural architecture search on target task and hardware. *arXiv preprint arXiv:1812.00332*.
- Cai, Z.; He, X.; Sun, J.; and Vasconcelos, N. 2017. Deep learning with low precision by half-wave gaussian quantization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 5918–5926.
- Chen, L.-C.; Papandreou, G.; Schroff, F.; and Adam, H. 2017. Rethinking atrous convolution for semantic image segmentation. *arXiv preprint arXiv:1706.05587*.
- Chen, X.; Xie, L.; Wu, J.; and Tian, Q. 2019. Progressive Differentiable Architecture Search: Bridging the Depth Gap between Search and Evaluation. *arXiv preprint arXiv:1904.12760*.
- Choi, J.; Wang, Z.; Venkataramani, S.; Chuang, P. I.-J.; Srinivasan, V.; and Gopalakrishnan, K. 2018. Pact: Parameterized clipping activation for quantized neural networks. *arXiv preprint arXiv:1805.06085*.
- Courbariaux, M.; Hubara, I.; Soudry, D.; El-Yaniv, R.; and Bengio, Y. 2016. Binarized neural networks: Training deep neural networks with weights and activations constrained to +1 or -1. *arXiv preprint arXiv:1602.02830*.
- Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. ImageNet: A Large-Scale Hierarchical Image Database. In *CVPR09*.
- Ding, R.; Chin, T.-W.; Liu, Z.; and Marculescu, D. 2019. Regularizing activation distribution for training binarized deep networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 11408–11417.
- Dong, Z.; Yao, Z.; Gholami, A.; Mahoney, M.; and Keutzer, K. 2019. HAWQ: Hessian AWARE Quantization of Neural Networks with Mixed-Precision. *arXiv preprint arXiv:1905.03696*.
- Everingham, M.; Eslami, S. A.; Van Gool, L.; Williams, C. K.; Winn, J.; and Zisserman, A. 2015. The pascal visual object classes challenge: A retrospective. *International journal of computer vision* 111(1): 98–136.
- Figurnov, M.; Collins, M. D.; Zhu, Y.; Zhang, L.; Huang, J.; Vetrov, D.; and Salakhutdinov, R. 2017. Spatially adaptive computation time for residual networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 1039–1048.
- Goodfellow, I.; Shlens, J.; and Szegedy, C. 2015. Explaining and Harnessing Adversarial Examples. In *International Conference on Learning Representations*. URL <http://arxiv.org/abs/1412.6572>.
- Hariharan, B.; Arbelaez, P.; Bourdev, L.; Maji, S.; and Malik, J. 2011. Semantic Contours from Inverse Detectors. In *International Conference on Computer Vision (ICCV)*.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2015. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *Proceedings of the IEEE international conference on computer vision*, 1026–1034.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.
- Hinton, G.; Vinyals, O.; and Dean, J. 2015. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*.
- Howard, A.; Sandler, M.; Chu, G.; Chen, L.-C.; Chen, B.; Tan, M.; Wang, W.; Zhu, Y.; Pang, R.; Vasudevan, V.; et al. 2019. Searching for mobilenetv3. *arXiv preprint arXiv:1905.02244*.
- Ioffe, S.; and Szegedy, C. 2015. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167*.
- 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 *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 4350–4359.
- Kingma, D. P.; and Ba, J. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Krizhevsky, A.; Hinton, G.; et al. 2009. Learning multiple layers of features from tiny images. Technical report, Cite-seer.
- Krizhevsky, A.; Sutskever, I.; and Hinton, G. E. 2012. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, 1097–1105.
- Liu, C.; Zoph, B.; Neumann, M.; Shlens, J.; Hua, W.; Li, L.-J.; Fei-Fei, L.; Yuille, A.; Huang, J.; and Murphy, K. 2018a. Progressive neural architecture search. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 19–34.
- Liu, Z.; Wu, B.; Luo, W.; Yang, X.; Liu, W.; and Cheng, K.-T. 2018b. Bi-real net: Enhancing the performance of 1-bit cnns with improved representational capability and advanced training algorithm. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 722–737.
- Nagel, M.; van Baalen, M.; Blankevoort, T.; and Welling, M. 2019. Data-Free Quantization through Weight Equalization and Bias Correction. *arXiv preprint arXiv:1906.04721*.
- Netzer, Y.; Wang, T.; Coates, A.; Bissacco, A.; Wu, B.; and Ng, A. Y. 2011. Reading digits in natural images with unsupervised feature learning. In *NIPS Workshop on Deep Learning and Unsupervised Feature Learning*.
- Paszke, A.; Gross, S.; Chintala, S.; Chanan, G.; Yang, E.; DeVito, Z.; Lin, Z.; Desmaison, A.; Antiga, L.; and Lerer, A. 2017. Automatic Differentiation in PyTorch. In *NIPS Autodiff Workshop*.

- Phan, H.; He, Y.; Savvides, M.; Shen, Z.; et al. 2020. Mobinet: A mobile binary network for image classification. In *The IEEE Winter Conference on Applications of Computer Vision*, 3453–3462.
- Rastegari, M.; Ordonez, V.; Redmon, J.; and Farhadi, A. 2016. Xnor-net: Imagenet classification using binary convolutional neural networks. In *European Conference on Computer Vision*, 525–542. Springer.
- Sandler, M.; Howard, A.; Zhu, M.; Zhmoginov, A.; and Chen, L.-C. 2018. Mobilenetv2: Inverted residuals and linear bottlenecks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 4510–4520.
- Tan, M.; and Le, Q. V. 2019. EfficientNet: Rethinking Model Scaling for Convolutional Neural Networks. *arXiv preprint arXiv:1905.11946*.
- Tang, W.; Hua, G.; and Wang, L. 2017. How to train a compact binary neural network with high accuracy? In *Thirty-First AAAI Conference on Artificial Intelligence*.
- Teerapittayanon, S.; McDanel, B.; and Kung, H.-T. 2016. Branchynet: Fast inference via early exiting from deep neural networks. In *2016 23rd International Conference on Pattern Recognition (ICPR)*, 2464–2469. IEEE.
- Veit, A.; and Belongie, S. 2018. Convolutional networks with adaptive inference graphs. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 3–18.
- Wang, K.; Liu, Z.; Lin, Y.; Lin, J.; and Han, S. 2019. HAQ: Hardware-Aware Automated Quantization with Mixed Precision. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 8612–8620.
- Wu, Z.; Nagarajan, T.; Kumar, A.; Rennie, S.; Davis, L. S.; Grauman, K.; and Feris, R. 2018. Blockdrop: Dynamic inference paths in residual networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 8817–8826.
- Yu, J.; and Huang, T. 2019. Universally slimmable networks and improved training techniques. *arXiv preprint arXiv:1903.05134*.
- Yu, J.; Yang, L.; Xu, N.; Yang, J.; and Huang, T. 2018. Slimmable neural networks. *arXiv preprint arXiv:1812.08928*.
- Zhang, D.; Yang, J.; Ye, D.; and Hua, G. 2018. Lq-nets: Learned quantization for highly accurate and compact deep neural networks. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 365–382.
- Zhao, X.; Li, H.; Shen, X.; Liang, X.; and Wu, Y. 2018. A modulation module for multi-task learning with applications in image retrieval. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 401–416.
- 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*.
- Zhuang, B.; Shen, C.; Tan, M.; Liu, L.; and Reid, I. 2018. Towards effective low-bitwidth convolutional neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 7920–7928.