

FracBits: Mixed Precision Quantization via Fractional Bit-Widths

Linjie Yang,^{1*} Qing Jin^{2*}

¹ByteDance Inc.

²Northeastern University

linjie.yang@bytedance.com jinqingking@gmail.com

Abstract

Model quantization helps to reduce model size and latency of deep neural networks. Mixed precision quantization is favorable with customized hardware supporting arithmetic operations at multiple bit-widths to achieve maximum efficiency. We propose a novel learning-based algorithm to derive mixed precision models end-to-end under target computation constraints and model sizes. During the optimization, the bit-width of each layer / kernel in the model is at a fractional status of two consecutive bit-widths which can be adjusted gradually. With a differentiable regularization term, the resource constraints can be met during the quantization-aware training which results in an optimized mixed precision model. Our final models achieve comparable or better performance than previous quantization methods with mixed precision on MobilenetV1/V2, ResNet18 under different resource constraints on ImageNet dataset.

Introduction

Neural network quantization (Choi et al. 2018; Elthakeb et al. 2018; Jin, Yang, and Liao 2019b; Lou et al. 2020; Rastegari et al. 2016; Uhlich et al. 2020; Wang et al. 2019; Wu et al. 2018; Zhou et al. 2017, 2016) has attracted large amount of attention due to the resource and latency constraints in real applications. Recent progress on neural network quantization has shown that the performance of quantized models can be as good as full precision models under moderate target bit-width such as 4 bits (Jin, Yang, and Liao 2019b). Besides, customized hardware can be configured to support multiple bit-widths for neural networks (Jin, Yang, and Liao 2019a). In order to fully exploit the power of model quantization, mixed precision quantization strategies are proposed to strike a better balance between computation cost and model accuracy. With more flexibility to distribute the computation budgets across layers (Elthakeb et al. 2018; Jin, Yang, and Liao 2019b; Wu et al. 2018), or even weight kernels (Lou et al. 2020), the quantized models with mixed precision usually achieve favorable performance than the ones with uniform precision.

Current approaches for mixed precision quantization usually borrow ideas from neural architecture search (NAS) lit-

Methods	HAQ	ReLeQ	AutoQ	DNAS	US	DQ	Ours
differentiable	✗	✗	✗	✓	✗	✓	✓
one-shot	✗	✗	✗	✗	✗	✓	✓
layer-wise	✓	✓	✓	✓	✓	✓	✓
kernel-wise	✗	✗	✓	✗	✗	✗	✓

Table 1: A comparison of our approach and previous mixed quantization algorithms. Our method FracBits achieves one-shot differentiable search and supports both layer-wise and kernel-wise quantization.

erature. Suppose we have a neural network with each convolution layer consisting of N branches where each branch is the quantized convolution with different bit-width. Finding the best configuration for a mixed precision model can be achieved by preserving a single branch for each convolution layer and pruning all other branches, which is conceptually equivalent to some recent NAS algorithms that aim at searching sub-networks from a supergraph (Cai, Zhu, and Han 2018; Pham et al. 2018; Wu et al. 2019; Xie et al. 2018). ENAS (Pham et al. 2018) and SNAS (Xie et al. 2018) employ reinforcement learning (RL) to learn a policy to sample network blocks from a supergraph. ReLeQ (Elthakeb et al. 2018) and HAQ (Wang et al. 2019) follow this footprint and employ reinforcement learning to choose layer-wise bit-width configurations for a neural network. AutoQ (Lou et al. 2020) further optimizes bit-width of each convolution kernel using a hierarchical RL strategy. ProxylessNAS (Cai, Zhu, and Han 2018) and FBNet (Wu et al. 2019) adopt a path sampling method to jointly learn model weights and importance scores of each operation in the supergraph. DNAS (Wu et al. 2018) directly reuses this path sampling methods and adds a regularization term proportional to the computation cost or model size, in order to discover mixed precision models with a good trade-off between computational resources and accuracy. Uniform Sampling (US) (Guo et al. 2019) is a similar method which uses uniform sampling to sample subnetworks from the supergraph in training and then searches for pruned or quantized models using evolutionary algorithm.

However, previous approaches on mixed precision quantization mostly directly adopts NAS algorithms and do not leverage specific properties of quantized models. Different from NAS and model pruning, the quantitative differ-

*Equal Contribution

ence of weights and activations with similar bits is small. For example, choosing 4 or 5 bits for one weight matrix only generates around 7.4% difference in value, assuming weights are uniformly distributed on $[0, 1]$ with linear quantization scheme. Thus the transition from one bit to its neighboring bits can be considered as a differentiable operation with appropriate parameterization. Recently, DQ (Uhlich et al. 2020) utilizes the Straight-Through Estimation (Bengio, Léonard, and Courville 2013) to facilitate differentiable bit-switching by treating bit-width of each layer as continuous parameters. Here, we propose a new approach to treat the bit-widths as continuous values by interpolating quantized weights or activation values of its two neighboring bit-widths. Such an approach facilitates an efficient one-shot differentiable optimization procedure of mixed precision quantization. By allocating differentiable bit-widths to layers or kernels, it can enable both layer-wise and kernel-wise quantization. A high-level comparison of our methods and previous mixed precision methods is shown in Table 1.

In summary, our contribution of this work is threefold.

- We propose a fractional bit-widths formulation that creates a smooth transition between neighboring quantized bits of weights and activations, facilitating differentiable search in layer-wise or kernel-wise precision dimension.
- Our mixed precision quantization algorithm only needs one-shot training of the network, greatly reduces exploration cost for resource restrained tasks.
- Our simple and straight-forward formulation is ready to be used for different quantization schemes. We showed superior performance than uniform precision approaches and previous mixed precision approaches on a wide range of models with different quantization schemes.

Related Work

Quantized Neural Networks Previous quantization techniques can be categorized into two types. The first type named post-training quantization directly quantizes weights and activations of a pretrained full-precision model into lower bit (Krishnamoorthi 2018; Nagel et al. 2019). Another type of techniques named quantization-aware training is proposed to incorporate quantization into training stage. Early studies in this direction employ a single precision for the whole neural network. For example, DoReFa (Zhou et al. 2016) proposes to transform the unbounded weights into a finite interval to reduce undesired quantization error introduced by infrequent large outliers. PACT (Choi et al. 2018) investigates the effect of clipping activations from different layers, finding the layer-dependence of the optimal clipping-levels. SAT (Jin, Yang, and Liao 2019b) investigates the gradient scales in training with quantized weights, and further improves model performance by adjusting weight scales. As another direction, some work assigns different bit-widths to different layers or kernels, enabling more flexible computation budget allocation. The first attempts employ reinforcement learning technique with rewards from estimated memory and computational cost by formulas (Elthakeb et al. 2018) or simulators (Wang et al. 2019). AutoQ (Lou et al. 2020) modifies the training procedure into a hierarchical

strategy, resulting in fine-grained kernel-wise quantization. However, these RL strategies needs to sample and train a large number of model variants which is very resource-demanding. DNAS (Wu et al. 2018) resorts to a differentiable strategy by constructing a supernet with each layer comprised by a linear combination of outputs from different bit-widths. However, due to the discrepancy between the search process and final configuration, it still needs to re-train the discovered model candidates. To further improve the searching efficiency, we propose a one-shot differentiable search method with fractional bit-widths. Due to the smooth transition between fractional bit-width and final integer bit-width, our method embeds the bit-width searching and model finetuning stages in a single pass of model training. Meanwhile, our technique is also orthogonal to Uniform Sample (US) (Guo et al. 2019) which trains a supernet by uniform sampling and searches good sub-architectures with evolutionary algorithm.

Network Pruning Network pruning is an orthogonal approach to speed up inference of neural networks to quantization. Early work (Han, Mao, and Dally 2015) compresses bulky models by learning connection together with weights, which produces unstructured connection in the final network. Later, structured compression by kernel-wise (Luo, Wu, and Lin 2017) or channel-wise (Gordon et al. 2018; He, Zhang, and Sun 2017; Liu et al. 2017; Ye et al. 2018) pruning is proposed, where the learned architecture is more friendly with acceleration on modern hardware. As an example, (Liu et al. 2017) identifies and prunes insignificant channels in each layer by penalizing on the scaling factor of the batch normalization layer. More recently, NAS algorithms are leveraged to guide network pruning. (Yu and Huang 2019) presents a one-shot searching algorithm by greedily slimming a pretrained slimmable neural network (Yu et al. 2018). (Mei et al. 2019) proposes a one-shot resource-aware searching algorithm using FLOPs as an L1 regularization term on the scaling factor of the batch normalization layer. We adopt a similar strategy to use BitOPs and model sizes as L1 regularization which are computed based on the trainable fractional bit-widths in our framework.

Mixed Precision Quantization

In this section, we will introduce our proposed method for mixed precision quantization. Our one-shot training pipeline involves two steps: bit-width searching and finetuning. We first introduce the implementation of fractional bit-width, and integration of the resource constraint in the searching process. After that, we introduce implementation of kernel-wise mixed precision quantization.

Searching with Fractional Bit-widths

In order to learn bit-widths dynamically in one-shot training, it is necessary to make them differentiable and define their derivative accordingly. To this end, we first examine a generic operation $f_k(x)$ that quantizes a value x to k -bit. Typically, $f_k(x)$ is well-defined only for positive integer values of k . To generalize bit-width to an arbitrary positive real

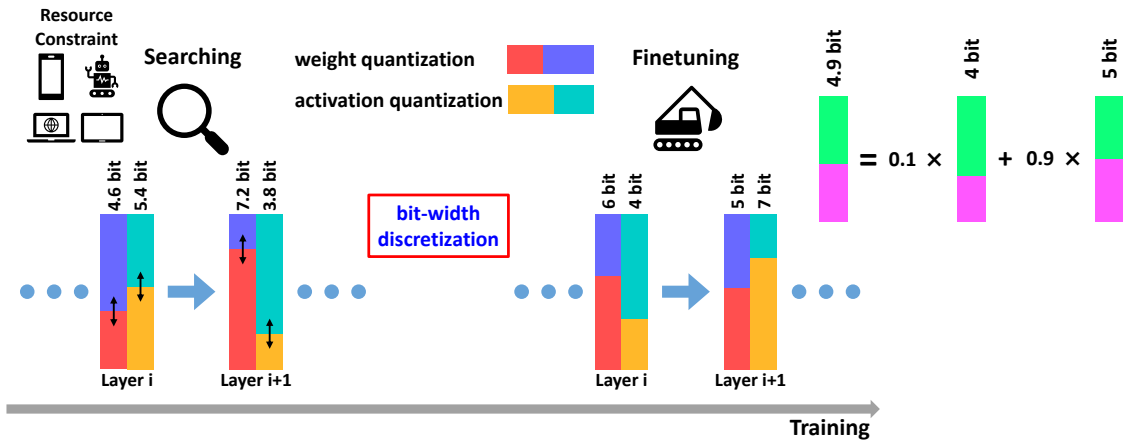


Figure 1: Our differentiable bit-width searching method consists of two stages: searching with fractional bit-width and finetuning with mixed bit-width quantization.

number λ , we apply first-order expansion around one of its nearby integer, and approximate the derivative at this integer by the slope of the segment joining the two adjacent grid points neighboring λ . Such a linear interpolation reads

$$f_\lambda(x) \approx f_{\lfloor \lambda \rfloor}(x) + (\lambda - \lfloor \lambda \rfloor)(f_{\lceil \lambda \rceil}(x) - f_{\lfloor \lambda \rfloor}(x)) \quad (1)$$

where $\lfloor \cdot \rfloor$ and $\lceil \cdot \rceil$ denote the floor and ceiling function, respectively. In other words, we can approximate an operation with a fractional bit-width by a linear combination of two operations with integer bit-widths, thus naturally achieving differentiability on it and making it learnable through typical gradient-based optimization, such as SGD. Note that the approximation in (1) turns into a strict equality if the original operation $f_k(x)$ is linear in k or if λ takes an integer value. The basic idea is illustrated in Fig. 1. In (1), the two rounding functions floor and ceiling on bit-width has vanishing gradient with respect to the argument, and thus the partial derivative of (1) with respect to λ is given by

$$\frac{\partial}{\partial \lambda} f_\lambda(x) = f_{\lceil \lambda \rceil}(x) - f_{\lfloor \lambda \rfloor}(x) \quad (2)$$

The difference of such an linear interpolation scheme compared to the widely-adopted straight through estimation (STE) (Bengio, Léonard, and Courville 2013) is that it uses soft bit-widths in both forward and backward propagation, rather than hard bit-widths in forward and soft bit-widths in back-propagation, as adopted by (Uhlich et al. 2020). In this way, the computed gradient reflects the true direction that the network parameters need to evolve along which results in better convergence.

Throughout we will adopt the DoReFa scheme for weight quantization, and the PACT scheme for activation quantization. The quantization function for both is the same, defined as

$$q_k(x) = \frac{1}{a} \lfloor ax \rfloor \quad (3)$$

where $x \in [0, 1]$, $\lfloor \cdot \rfloor$ indicates rounding to the nearest integer, and a equals $2^k - 1$ where k is the quantization bit-width. Thus, for both quantization, we have $f_k(x) = q_k(x)$ for integer bit-widths, and quantization with fractional bit-widths is

implemented with Eq. (1). The weight quantization is given by $Q_W = 2q_{\lambda_w}(\widetilde{W}) - 1$, where \widetilde{W} is the transformed weight clamped to the interval $[0, 1]$; activation quantization is given by $Q_X = \alpha q_{\lambda_a}(\frac{X}{\alpha})$, where α is a learnable parameter and X is the original activation clipped at α . λ_w and λ_a are the learnable fractional bit-widths for weight and activation, respectively. Also, it is possible to privatize bit-width to each kernel, enabling kernel-wise mixed precision quantization, as discussed later.

During the earlier searching stage, the precision assigned to each layer or each kernel is still undetermined, and we want to find the optimal bit-width structure through training. By initializing each bit-width with some arbitrary value, we can use (1) to quantize weights and activations in the model to fractional bit-widths. Meanwhile, this allows us to assign different bit-widths to different layers or even kernels, as well as to furnish separate precision for weight and activation quantization. During the training process, the model gradually converges to an optimal bit-width for both weight and activation corresponding to each unit, enabling quantization with mixed precision.

Resource Constraint as Penalty Loss

Restricting storage or computation cost is essential for model quantization, as the original purpose of quantization is to save resource consumption when deploying bulky model on portable devices or embedded systems. To this end, previous work resort to constraining on different metrics during the optimization procedure, including memory footprints (Uhlich et al. 2020), model size (Uhlich et al. 2020; Wang et al. 2019), BitOPs (Guo et al. 2019; Wu et al. 2018) and even estimated latency or energy (Lou et al. 2020; Wang et al. 2019). Here, we focus on model size in bits (Bytes) for weight-only quantization, and the number of BitOPs for quantization on both weight and activation, as they can be directly calculated from assigned bit-widths. Note latency and energy consumption (Lou et al. 2020; Wang et al. 2019) may seem to be more practical measures for real applications. However, we argue that BitOPs

can also be a good metric since it is solely determined by the model itself rather than different configurations of hardwares, simulators and compilers, which guarantees fair comparison between different approaches and advocates reproducible research.

Weight-only quantization targets at shrinking the model size, while floating point operation is still needed during inference. Model size are usually expressed in terms of the required number of bits to store weights (and bias) in the model. For a weight w of k_w -bit, the size is simply k_w . The generalized model size for a fractional bit-width λ_w is thus λ_w . The size of the whole model can be obtained by summing over all weights in the model. Note that the bit-width can be shared among all weights in the whole layer or along each kernel (as discussed later), corresponding to layer-wise or kernel-wise quantization, respectively. For example, for a typical 2D convolution layer (without grouping) sharing the same fractional bit-width λ_w among all weights, the size is given by $\lambda_w c_{in} c_{out} k_x k_y$, where c_{in} is the number of input channels, c_{out} is the number of output channels, and k_x and k_y represent the horizontal and vertical kernel sizes, respectively.

Quantization on both weights and activations can effectively decrease computation cost for real application, which can be measured with number of BitOPs involved in multiplications. Suppose a weight value w and an activation value a involved in multiplication are quantized to k_w -bit and k_a -bit, respectively. The number of BitOPs for such a multiplication is

$$\text{comp}_{wa} = k_w k_a \quad (4)$$

This expression is bi-linear in k_w and k_a , which means that for fractional bit-widths λ_w and λ_a , (1) leads to

$$\text{size}_w = \lambda_w \quad (5)$$

$$\text{comp}_{wa} = \lambda_w \lambda_a \quad (6)$$

The total computation cost of the model is the sum over all weights and activations. As for the example of 2D convolution layer, if all weights share the same fractional bit-width λ_w and all input activations share the same fractional bit-width λ_a , the number of BitOPs is given by $\lambda_w \lambda_a c_{in} c_{out} k_x k_y o_x o_y$, where o_x and o_y represents the horizontal and vertical sizes of the output features, respectively.

Targeting prescribed objective With constraints defined properly, we are able to penalize on them to enable constraint-aware optimization. Here, we directly define the penalty term as the L1 difference from some target constraint value by

$$\mathcal{L}^{\text{size}} = \left| \sum_w \text{size}_w - \text{size}_t \right| \quad (7a)$$

$$\mathcal{L}^{\text{comp}} = \left| \sum_{wa} \text{comp}_{wa} - \text{comp}_t \right| \quad (7b)$$

where size_t and comp_t denote target constraints for model size and computation cost, respectively. The sum is taken over all weights in the model for model size constrained optimization, and is taken over all weights and all activations for computation cost constrained case. Following the convention adopted in most literature, for both constraints we

only take into account those contributed by convolution and fully-connected layers.

Adding the penalty term to the original loss (such as cross entropy for classification task) with a coefficient κ , we arrive at the total loss for optimization

$$\mathcal{L}^{\text{total}} = \mathcal{L}^{\text{cls}} + \kappa \cdot \mathcal{L}^{\text{size}} \quad (8a)$$

$$\mathcal{L}^{\text{total}} = \mathcal{L}^{\text{cls}} + \kappa \cdot \mathcal{L}^{\text{comp}} \quad (8b)$$

It should be noted that the value of κ depends on the unit of constraints. Throughout the paper, we measure model size in terms of MB (megabytes) and computation cost in terms of GBitOPs (billion of BitOPs). In this way, the desired resource constraint can be reached in the joint optimization of model parameters and bit-widths. Note that the recent concurrent work (Nikolić et al. 2020) adopts a similar approach for mixed precision quantization with L1 regularization on bit-widths for weights and activations, while here we explicitly define the loss as a function of computational cost in BitOPs or model size in Bytes and incorporate the target constraint into the loss directly.

Finetuning with Mixed Precision

After searching, the bit-widths of the model are still continuous values. We discretize the bit-widths with a threshold value to make the resulted network meeting the resource constraints. Specifically, we use binary search to find a threshold for both weight and activation bit-widths, which makes the resource cost of the model to be within 1% deviation of the target constraint. This way, each layer or each kernel has its individual bit-widths for weights and activations learned in the previous stage, and the training enters the finetuning stage to only update model weights. The ratio between training epochs allocated to searching and finetuning is a hyper-parameter that can be freely specified. In practice, we assign 80% of training epochs to searching and 20% to finetuning. Here we want to emphasize that the combination of searching and finetuning constitutes the whole training procedure, and the total number of epochs of the two stages is the same as a traditional quantization-aware training procedure. Also, the whole procedure for updating learned parameters and scheduling hyper-parameter (learning rate, weight decay, etc.) is smooth, and do not need any re-initialization for the finetuning. Thus, our training method is one-shot, without extra retraining steps.

Kernel-wise Mixed Precision Quantization

As mentioned above, our algorithm is not restricted to layer-wise quantization, but also supports kernel-wise quantization. Here, one kernel means weight parameters associated with a convolution filter to produce a single-channel feature map. Weight kernels in a convolution layer are assigned with different bit-width parameters λ_{w_i} , where i is the index of the weight kernel. For each convolution operation of one weight kernel with the input tensor, the input tensor can also be assigned with different bit-widths. However, quantizing the input tensor with different bit-widths for different weight kernels requires large computation overhead. Here we assign the same bit-width λ_a on the input

tensor for computation with all the weight kernels. Note that (Lou et al. 2020) adopted the same strategy for kernel-wise quantization. For a 2D convolution layer, the number of BitOPs associated with the fractional bit-width is given by $\sum_i \lambda_{w_i} \lambda_{a_{in}} k_x k_y o_x o_y$, and model size can be represented by $\sum_i \lambda_{w_i} c_{in} k_x k_y$.

Experiments

In this section, we conduct quantitative experiments using FracBits and compare it with previous quantization approaches including uniform quantization algorithms PACT (Choi et al. 2018), LQNet (Zhang et al. 2018), SAT (Jin, Yang, and Liao 2019b) and mixed precision quantization algorithms HAQ (Wang et al. 2019), AutoQ (Lou et al. 2020), DNAS (Wu et al. 2018), US (Guo et al. 2019), DQ (Uhlich et al. 2020). We first compare our method with previous approaches on layer-wise mixed precision quantization. Then we compare our method with a previous kernel-wise mixed precision method AutoQ on kernel-wise precision search.

Implementation Details

We build our algorithms based on a recent quantization algorithm SAT (Jin, Yang, and Liao 2019b), which is an improved version of PACT algorithm (Choi et al. 2018). PACT jointly learns quantized weights and activations where weights are quantized using the DoReFa scheme (Zhou et al. 2016), while SAT modifies PACT with gradient calibration and scale adjusting. κ is a critical parameter for the proper convergence of the network towards required resource constraints. Models under mild or aggressive constraints may couple with different values of κ . Different types of resource constraints (computational cost and model size) have different scales and requires different scales of the regularization term. However, in our experiments, we find our algorithm is not very sensitive to values of κ . We set κ to 0.1 for all computation cost constrained experiments, and 1 for all model size constrained experiments. We also find it beneficial to initialize the model at some point close to the target resource constraint, facilitating more exploration close to the target model spaces. We set the initial value of κ to $b_t + 0.5$ in each layer for all experiments, where b_t is the bit-width achieving similar resource constraints in the corresponding uniformly quantized model. For all channel-wise quantization experiments with both weights and activations quantized, we set the candidate bit-widths to be 2-8. For all other experiments including weight-only quantization and kernel-wise quantization, we set the candidate bit-widths to be 1-8. Since the first and the last layers in a neural network have crucial impact on the performance of the model, we fix the bit-width of the first and last layer to 8 bit following (Jin, Yang, and Liao 2019b).

For all experiments, we use cosine learning rate scheduler without restart. Learning rate is initially set to 0.05 and updated every iteration for totally 150 epochs. We use SGD optimizer with a momentum weight of 0.9 without damping, and weight decay of 4×10^{-5} . The batch size is set to 2048 for all models. The warmup strategy suggested in (Goyal

	bit-width	3			4		
	method	top-1	top-5	bitops	top-1	top-5	bitops
MBNetV1	HAQ	-	-	-	67.4	87.9	-
	SAT	67.1	87.1	5.73	71.3	89.9	9.64
	FracBits-SAT	68.7	88.2	5.78	71.4	90.0	9.63
MBNetV2	HAQ	-	-	-	67	87.3	-
	AutoQ	-	-	-	69	89.4	-
	DQ	-	-	-	69.7	-	-
	SAT	67.2	87.3	3.32	70.8	89.7	5.35
	FracBits-SAT	67.8	87.6	3.33	71.3	90.1	5.35

Table 2: Comparison of computation cost constrained layer-wise quantization of our method and previous approaches on ImageNet with MobileNet V1/V2. Note that accuracies are in % and bitops are in B (billion).

et al. 2017) is also adopted by linearly increasing the learning rate every iteration to $batchsize/256 \times 0.05$ for the first five epochs before using the cosine annealing scheduler. Bit-width search is conducted in the first 120 epochs after the warmup stage. At the 121th epoch, all fractional bit-width will be discretized to integer bits, and the network will be further finetuned for the rest 30 epochs. We do not observe any glitch in the training loss in this discretization process, potentially due to the insignificant difference in quantized values of two neighboring bit-widths.

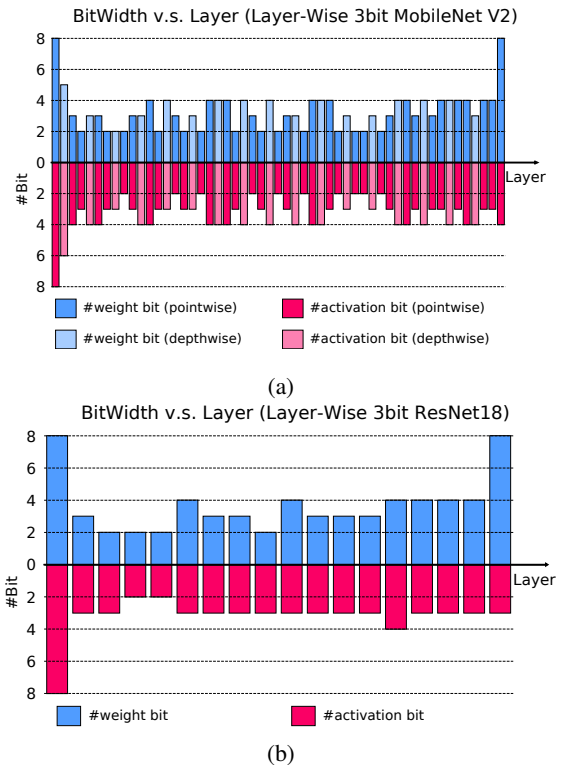


Figure 2: Layer-wise mixed precision quantization for 3 bit MobileNet V2 (a) and ResNet18 (b).

bit-width	3			4			FP
	top-1	Δ_{acc}	bitops	top-1	Δ_{acc}	bitops	top-1
PACT (Choi et al. 2018)	68.3	-1.9	22.83	69.2	-1.0	34.70	70.2
LQNet (Zhang et al. 2018)	68.2	-2.1	22.83	69.3	-1.0	34.70	70.3
DNAS (Wu et al. 2018)	68.7	-2.3	24.34*	70.6	-0.4	35.17*	71.0
DQ (Uhlich et al. 2020)	-	-	-	70.1	-0.2	-	70.3
AutoQ (Lou et al. 2020)	-	-	-	68.2	-1.7	-	69.9
US (Guo et al. 2019)	69.4	-1.5	22.11*	70.5	-0.4	33.74*	70.9
FracBits-PACT	69.1	-1.1	22.70	69.7	-0.5	34.73	70.2
SAT (Jin, Yang, and Liao 2019b)	69.3	-0.9	22.83	70.3	0.1	34.70	70.2
FracBits-SAT	69.4	-0.8	22.93	70.6	0.4	34.70	70.2

Table 3: Comparison of computation cost constrained layer-wise quantization of our method and previous approaches on ImageNet with ResNet18. Note bitops of US (Guo et al. 2019) and DNAS (Wu et al. 2018) does not include first and last layer in their papers, and US shows different bitops numbers from ours. We give an estimation of their bitops based on the difference with uniformly quantized models. Note that accuracies are in % and bitops are in B (billion).

Quantization with Layer-wise Precision

We compare FracBits with previous quantization algorithms on layer-wise precision search. We conducted experiments on MobileNet V1/V2 and ResNet18. Since FracBits can be used for both computation cost constrained and model size constrained bit-width search, we conduct experiments on both settings to validate the effectiveness of our approach.

Table 2 shows experiment results of layer-wise computation cost constrained quantization on MobileNet V1/V2. Previous methods HAQ (Wang et al. 2019) and AutoQ (Lou et al. 2020) use PACT as quantization scheme, while DQ uses a similar scheme to PACT with learnable clipping bounds. Derived from PACT, SAT is a strong uniform quantization baseline which already outperforms all previous mixed precision methods. For example, it already achieves 71.3% on 4-bit MobileNet V1 and 70.8% on 4-bit MobileNet V2, almost closing the gap between full precision models (71.7% for MobileNet V1 and 71.8% for MobileNet V2) and quantized ones. We believe that validating the effectiveness of our FracBits algorithm based on SAT is helpful towards seeking the limit of mixed precision quantization algorithms. We find that FracBits-SAT achieves slightly better performance compared to SAT on 4-bit MobileNet V1/V2, and achieves significantly better result on 3-bit models, which proves its effectiveness on strong uniform quantization baselines. It has a 1.6% absolute gain on 3-bit MobileNet V1 and a 0.6% gain on MobileNet V2¹ under the same computation cost budget.

We show comparison with more algorithms on ResNet18, as listed in Table 3. Here we compare with uniform precision approaches PACT, LQNet and mixed precision approaches DNAS, DQ, AutoQ, and US. Except DQ, all mixed precision approaches use PACT as quantization scheme. Since all methods report different accuracies for full precision (FP) models, we also add the top-1 accuracy of FP models reported in corresponding papers and report the relative ac-

¹In MobileNet V2, some convolution layers are not followed by ReLU activation. Here we use double-sided quantization for outputs of these layers, meaning that they are clipped into an interval of $[-\alpha, \alpha]$. Our re-implemented SAT also adopts such double-side clipping.

curacy drop for each method. Here we also include results of directly applying our method on PACT scheme, and denote this as FracBits-PACT. Comparing absolute accuracy, FracBits-PACT achieves comparable performance as state-of-the-art mixed precision methods. Note DNAS uses several tricks in training to boost performance, thus its result is not directly comparable to others. Comparing relative accuracy drop, our method achieves least performance drop on 3-bit ResNet18. Enhanced by SAT quantization method, FracBits-SAT further improves over SAT baseline and achieves only 0.8% accuracy drop on 3-bit ResNet18 and even a 0.4% performance gain on 4-bit ResNet18. Note DQ and our method are one-shot differentiable method which only need one pass of training to obtain the final model, and are much more efficient than the other mixed quantization approaches (DNAS, AutoQ, US).

To have a more intuitive understanding of the learned bit-width structure from our algorithm, we plot the bit-widths from different layers for 3-bit MobileNet V2 and ResNet18, as shown in Fig. 2. We find that models for mixed quantization constrained on computational cost generally uses more bit-width on the late stage of the network, potentially due to the larger computation cost of early layers than later layers. Also, in MobileNet V2, depth-wise convolutions result in more bit-width than point-wise convolutions due to their low computation cost.

For model size constrained quantization, we show comparison with previous methods Deep Compression (Han, Mao, and Dally 2015), HAQ and uniform quantization approach SAT in Table 4. Our FracBits-SAT outperforms mixed precision methods HAQ and strong uniform quantization baseline SAT on all experimented bit-widths consistently. Note that FracBits has an over 3% absolute gain on top-1 accuracy over SAT on 2-bit MobileNet V1/V2. On the challenging 3-bit setting where quantized models already achieve similar performance as full precision ones, FracBits also outperforms SAT with a 0.6% margin on MobileNet V1 and a 0.8% gain on MobileNet V2 in top-1 accuracy.


	bit-width	2			3		
	method	top-1	top-5	size	top-1	top-5	size
MBNetV1	DeepComp	37.6	64.3	1.09	65.9	86.9	1.60
	HAQ	57.1	81.9	1.09	67.7	88.2	1.58
	SAT	66.3	86.8	1.83	70.7	89.5	2.22
	FracBits-SAT	69.7	88.9	1.81	71.3	90.0	2.23
MBNetV2	DeepComp	58.1	82.2	0.96	68.0	88.0	1.38
	HAQ	66.8	87.3	0.95	70.9	89.8	1.38
	SAT	66.8	87.2	1.83	71.1	89.9	2.11
	FracBits-SAT	69.9	89.3	1.84	71.9	90.4	2.12

Table 4: Comparison of model size constrained layer-wise quantization of our method and previous approaches on ImageNet with MobileNet V1/V2. Note that accuracies are in % and sizes are in MB. The difference in model size is due to that we use 8bit for the last fully-connected layer, following previous work (Choi et al. 2018; Jin, Yang, and Liao 2019b), while this bit-width is searched in (Wang et al. 2019).

	bit-width	3			4		
	method	top-1	top-5	bitops	top-1	top-5	bitops
MBNetV2	AutoQ	-	-	-	70.8	90.3	-
	FB-PACT-K	68.0	87.8	3.33	70.9	89.5	5.36
	SAT	67.2	87.3	3.32	70.8	89.7	5.35
	FB-SAT-K	68.2	87.9	3.35	71.6	90.0	5.33
ResNet18	AutoQ	-	-	-	69.8	88.4	-
	FB-PACT-K	69.0	88.3	23.01	69.9	88.8	34.70
	SAT	69.3	88.9	22.83	70.3	89.5	34.70
	FB-SAT-K	69.8	88.9	22.87	70.8	89.6	34.82


Table 5: Comparison of computation cost constrained kernel-wise quantization of our method and previous approaches on MobileNet V2 and ResNet18. Note that accuracies are in % and bitops are in B (billion). FB-PACT-K and FB-SAT-K denote FracBits for kernel-wise quantization with PACT and SAT quantization schemes, respectively.

Weight BitWidth Distribution v.s. Layer (Kernel-Wise 3bit MobileNet V2)



(a)

Weight BitWidth Distribution v.s. Layer (Kernel-Wise 3bit ResNet18)



(b)

Figure 3: Kernel-wise mixed precision quantization for 3 bit MobileNet V2 (a) and ResNet18 (b).

Quantization with Kernel-wise Precision

In this section, we experiment with quantization with kernel-wise precision. Among previous approaches, only AutoQ (Lou et al. 2020) has experiments on kernel-wise precision which we will compare with. In Table 5, we denote kernel-wise FracBits based on PACT and SAT as FB-PACT-K and FB-SAT-K, and compare them with AutoQ and uni-

form precision method SAT. FB-PACT-K achieves comparable results as AutoQ on MobileNet V2 and ResNet18, while being much more efficient than the RL based method which needs to train hundreds of model variants, thanks to the differentiable formulation. FB-SAT-K outperforms SAT significantly with 1.0% and 0.8% increase on top-1 accuracy on 3 and 4-bit MobileNet V2, respectively, and with 0.5% increase on top-1 accuracy on both 3 and 4-bit ResNet18. Compared to layer-wise precision models, FB-SAT-K outperforms FracBits-SAT by 0.4% and 0.3% on 3 and 4-bit MobileNet V2, respectively. It also outperforms layer-wise FracBits-SAT by 0.4% on 3-bit ResNet18, proving our kernel-wise quantization method can further improve over strong layer-wise mixed-precision models. Fig. 3 illustrates the bit-width distribution against layer indices for 3-bit MobileNet V2 and ResNet18. We can see that 3-bit MobileNet V2 assign low bits in the early layers and intermediate bottleneck layers, while 3-bit ResNet-18 assign low bits only in early layers. We believe that the point-wise convolutions in MobileNet V2 have much larger computation cost compared to depth-wise convolutions and thus they receive a larger resource penalty during optimization, leading to more compression by lower bit-widths.

Conclusion

We propose a new formulation named FracBits for mixed precision quantization. We formulate the bit-width of each layer or kernel with a continuous learnable parameter that can be instantiated by interpolating quantized parameters of two neighboring bit-widths. Our method facilitates differentiable optimization of layer-wise or kernel-wise bit-width in a single shot of training. With only a regularized term to penalize extra computational resource in the training process, our method is able to discover proper bit-width configurations for different models, outperforming previous mixed precision and uniform precision approaches. We believe our method will motivate research along low-precision neural networks, and low-cost computational models.

References

- 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* .
- 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* .
- Elthakeb, A. T.; Pilligundla, P.; Yazdanbakhsh, A.; Kinzer, S.; and Esmaeilzadeh, H. 2018. ReLeQ: A Reinforcement Learning Approach for Deep Quantization of Neural Networks. In *NuerIPS*.
- Gordon, A.; Eban, E.; Nachum, O.; Chen, B.; Wu, H.; Yang, T.-J.; and Choi, E. 2018. Morphnet: Fast & simple resource-constrained structure learning of deep networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 1586–1595.
- Goyal, P.; Dollár, P.; Girshick, R. B.; Noordhuis, P.; Wesolowski, L.; Kyrola, A.; Tulloch, A.; Jia, Y.; and He, K. 2017. Accurate, Large Minibatch SGD: Training ImageNet in 1 Hour. *CoRR* abs/1706.02677. URL <http://arxiv.org/abs/1706.02677>.
- Guo, Z.; Zhang, X.; Mu, H.; Heng, W.; Liu, Z.; Wei, Y.; and Sun, J. 2019. Single Path One-Shot Neural Architecture Search with Uniform Sampling. *CoRR* abs/1904.00420. URL <http://arxiv.org/abs/1904.00420>.
- Han, S.; Mao, H.; and Dally, W. J. 2015. Deep compression: Compressing deep neural networks with pruning, trained quantization and Huffman coding. *arXiv preprint arXiv:1510.00149* .
- He, Y.; Zhang, X.; and Sun, J. 2017. Channel pruning for accelerating very deep neural networks. In *Proceedings of the IEEE International Conference on Computer Vision*, 1389–1397.
- Jin, Q.; Yang, L.; and Liao, Z. 2019a. AdaBits: Neural Network Quantization with Adaptive Bit-Widths. *arXiv preprint arXiv:1912.09666* .
- Jin, Q.; Yang, L.; and Liao, Z. 2019b. Towards Efficient Training for Neural Network Quantization. *arXiv preprint arXiv:1912.10207* .
- Krishnamoorthi, R. 2018. Quantizing deep convolutional networks for efficient inference: A whitepaper. *arXiv preprint arXiv:1806.08342* .
- Liu, Z.; Li, J.; Shen, Z.; Huang, G.; Yan, S.; and Zhang, C. 2017. Learning efficient convolutional networks through network slimming. In *Proceedings of the IEEE International Conference on Computer Vision*, 2736–2744.
- Lou, Q.; Guo, F.; Kim, M.; Liu, L.; and Jiang, L. 2020. AutoQ: Automated Kernel-Wise Neural Network Quantization. In *International Conference on Learning Representations*. URL <https://openreview.net/forum?id=rygfn4twS>.
- Luo, J.-H.; Wu, J.; and Lin, W. 2017. Thinet: A filter level pruning method for deep neural network compression. In *Proceedings of the IEEE international conference on computer vision*, 5058–5066.
- Mei, J.; Li, Y.; Lian, X.; Jin, X.; Yang, L.; Yuille, A.; and Yang, J. 2019. AtomNAS: Fine-Grained End-to-End Neural Architecture Search. *arXiv preprint arXiv:1912.09640* .
- Nagel, M.; Baalen, M. v.; Blankevoort, T.; and Welling, M. 2019. Data-free quantization through weight equalization and bias correction. In *Proceedings of the IEEE International Conference on Computer Vision*, 1325–1334.
- Nikolić, M.; Hacene, G. B.; Bannon, C.; Lascorz, A. D.; Courbariaux, M.; Bengio, Y.; Gripon, V.; and Moshovos, A. 2020. BitPruning: Learning Bitlengths for Aggressive and Accurate Quantization. *arXiv preprint arXiv:2002.03090* .
- Pham, H.; Guan, M. Y.; Zoph, B.; Le, Q. V.; and Dean, J. 2018. Efficient neural architecture search via parameter sharing. *arXiv preprint arXiv:1802.03268* .
- 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.
- Uhlich, S.; Mauch, L.; Yoshiyama, K.; Cardinaux, F.; Garcia, J. A.; Tiedemann, S.; Kemp, T.; and Nakamura, A. 2020. MIXED PRECISION DNNs: ALL YOU NEED IS A GOOD PARAMETRIZATION. *ICLR* .
- 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, B.; Dai, X.; Zhang, P.; Wang, Y.; Sun, F.; Wu, Y.; Tian, Y.; Vajda, P.; Jia, Y.; and Keutzer, K. 2019. Fbnet: Hardware-aware efficient convnet design via differentiable neural architecture search. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 10734–10742.
- Wu, B.; Wang, Y.; Zhang, P.; Tian, Y.; Vajda, P.; and Keutzer, K. 2018. Mixed precision quantization of convnets via differentiable neural architecture search. *arXiv preprint arXiv:1812.00090* .
- Xie, S.; Zheng, H.; Liu, C.; and Lin, L. 2018. SNAS: stochastic neural architecture search. *arXiv preprint arXiv:1812.09926* .
- Ye, J.; Lu, X.; Lin, Z.; and Wang, J. Z. 2018. Rethinking the smaller-norm-less-informative assumption in channel pruning of convolution layers. *arXiv preprint arXiv:1802.00124* .
- Yu, J.; and Huang, T. S. 2019. Network Slimming by Slimmable Networks: Towards One-Shot Architecture Search for Channel Numbers. *CoRR* abs/1903.11728. URL <http://arxiv.org/abs/1903.11728>.
- 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.

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*.

Zhou, S.-C.; Wang, Y.-Z.; Wen, H.; He, Q.-Y.; and Zou, Y.-H. 2017. Balanced quantization: An effective and efficient approach to quantized neural networks. *Journal of Computer Science and Technology* 32(4): 667–682.